FY13 Implementation Plan

Vol. 2, Rev. 0.5

September 2012

LLNL-TR-584652

NA-ASC-125R-12-Vol.2-Rev.0.5-IP

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Advanced Simulation and Computing

FY13 IMPLEMENTATION PLAN
Volume 2, Rev. 0.5

September 27, 2012

Approved by:
Robert Meisner,
NNSA ASC Program Director 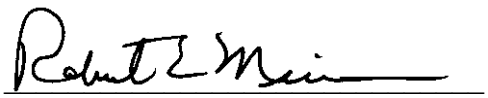 9/26/12
Signature Date

Michel McCoy,
LLNL ASC Executive *Concurrence Received* 9/20/12

Bill Archer,
LANL ASC Executive, acting *Concurrence Received* 9/20/12

Wendy Cieslak,
SNL ASC Executive *Concurrence Received* 9/20/12

ASC Focal Point	**IP Focal Point**	**IP Focal Point**
Robert Meisner	***Doug Wade***	***Thuc Hoang***
NA-114	NA-114	NA-114
Tele.: 202-586-0908	Tele.: 202-586-7507	Tele.: 202-586-7050
FAX: 202-586-0405	FAX: 202-586-0405	FAX: 202-586-0405
bob.meisner@nnsa.doe.gov	douglas.wade@hq.doe.gov	thuc.hoang@hq.doe.gov

Contents

I. Executive Summary

The Stockpile Stewardship Program (SSP) is a single, highly integrated technical program for maintaining the surety and reliability of the U.S. nuclear stockpile. The SSP uses nuclear test data, computational modeling and simulation, and experimental facilities to advance understanding of nuclear weapons. It includes stockpile surveillance, experimental research, development and engineering programs, and an appropriately scaled production capability to support stockpile requirements. This integrated national program requires the continued use of experimental facilities and programs, and the computational enhancements to support these programs.

The Advanced Simulation and Computing Program (ASC) is a cornerstone of the SSP, providing simulation capabilities and computational resources that support annual stockpile assessment and certification, study advanced nuclear weapons design and manufacturing processes, analyze accident scenarios and weapons aging, and provide the tools to enable stockpile Life Extension Programs (LEPs) and the resolution of Significant Finding Investigations (SFIs). This requires a balanced resource, including technical staff, hardware, simulation software, and computer science solutions.

In its first decade, the ASC strategy focused on demonstrating simulation capabilities of unprecedented scale in three spatial dimensions. In its second decade, ASC is now focused on increasing predictive capabilities in a three-dimensional (3D) simulation environment while maintaining support to the SSP. The program continues to improve its unique tools for solving progressively more difficult stockpile problems (sufficient resolution, dimensionality, and scientific details), quantify critical margins and uncertainties, and resolve increasingly difficult analyses needed for the SSP.

Moreover, ASC's business model is integrated and focused on requirements-driven products that address long-standing technical questions related to enhanced predictive capability in the simulation tools.

ASC must continue to meet three objectives:

- **Objective 1. Robust Tools.** Develop robust models, codes, and computational techniques to support stockpile needs such as refurbishments, SFIs, LEPs, annual assessments, and evolving future requirements.
- **Objective 2. Prediction through Simulation.** Deliver verified and validated physics and engineering tools to 1) enable simulations of nuclear weapons performance in a variety of operational environments and physical regimes, and 2) enable risk-informed decisions about the performance, safety, and reliability of the stockpile.
- **Objective 3. Balanced Operational Infrastructure.** Implement a balanced computing platform acquisition strategy and operational infrastructure to meet Directed Stockpile Work (DSW) and SSP needs for production and advanced/high-end simulation capabilities.

II. Introduction

Prior to the start of the nuclear testing moratorium in October 1992, the nuclear weapons stockpile was maintained through 1) underground nuclear testing and surveillance activities and 2) "modernization" (in other words, development of new weapons systems). A consequence of the nuclear test ban is that the safety, performance, and reliability of U.S. nuclear weapons must be ensured by other means for systems far beyond the lifetimes originally envisioned when the weapons were designed.

Following the start of the moratorium, the ASC Program now supports the National Nuclear Security Administration's (NNSA's) overarching goal of nuclear weapons stewardship: *"We continue to advance the Stockpile Stewardship Program to push the scientific and engineering boundaries needed to maintain our nuclear arsenal. It also means maintaining the basic science and engineering that is the foundation of the weapons program."*[1]

In 1996, ASCI—the Accelerated Strategic Computing Initiative—was established as an essential element of the SSP to provide nuclear weapons simulation and modeling capabilities.

In 2000, the NNSA was established to carry out the national security responsibilities of the Department of Energy (DOE), including maintenance of a safe, secure, and reliable stockpile of nuclear weapons and associated materials capabilities and technologies.

Shortly thereafter, in 2002, ASCI matured from an initiative to a recognized program and was renamed the ASC Program.

NNSA will carry out its responsibilities through the twenty-first century in accordance with the current Administration's vision and the Nuclear Posture Review (NPR) guidance. NNSA Administrator Thomas P. D'Agostino summarized[2] the NNSA objectives for the SSP as follows:

"Our fundamental national security responsibilities for the United States include:

- *Assuring the safety, security and reliability of the U.S. nuclear weapons stockpile while at the same time transforming the stockpile and the infrastructure that supports it;*
- *Reducing the threat posed by nuclear proliferation; and,*
- *Providing reliable and safe nuclear reactor propulsion systems for the U.S. Navy."*

"Throughout the past decade, the Stockpile Stewardship Program has proven its ability to successfully sustain the safety, security and reliability of the nuclear arsenal without

[1] *NNSA Strategic Planning Guidance for FY2010–2014*, April 2008, page 17.

[2] Testimony on the *FY 2008 National Defense Authorization Budget Request for the Department of Energy's NNSA* before the House Armed Services Subcommittee, March 20, 2007.

resorting to underground nuclear testing. The SSP also enables the U.S. to provide a credible strategic deterrent capability with a stockpile that is significantly smaller."

Additionally, key investment recommendations cited in the 2010 NPR are to:

- Strengthen the science, technology, and engineering base needed for conducting weapon system LEPs
- Mature advanced technologies to increase weapons surety
- Qualify weapon components and certify weapons without nuclear testing
- Provide annual stockpile assessments through weapons surveillance

This strategy includes developing and sustaining high-quality scientific staff, as well as supporting computational and experimental capabilities. [3]

The ASC Program plays a vital role in the NNSA infrastructure and its ability to respond to the NPR guidance. The program focuses on developing modern simulation tools that can provide insights into stockpile problems, providing tools with which designers and analysts can certify nuclear weapons, and guiding any necessary modifications in nuclear warheads and the underpinning manufacturing processes. Additionally, ASC is enhancing the predictive capability necessary to evaluate weapons effects, designing experiments, and ensuring test readiness.

ASC continues to improve its unique tools to solve progressively more difficult stockpile problems—with a focus on sufficient resolution, dimensionality, and scientific details—to enable Quantification of Margins and Uncertainties (QMU) and to resolve the increasingly difficult analyses needed for stockpile stewardship. The DSW program provides requirements for simulation. These requirements include planned LEPs, stockpile support activities and requirements for future capabilities. ASC's advanced, leading-edge technologies in high performance computing (HPC) and predictive capability meet DSW's short- and long-term needs, including the annual assessments and certifications, as well as SFIs.

The following section provides an overview of FY11 and upcoming ASC contributions to the SSP:[4]

ASC Contributions to the Stockpile Stewardship Program

In FY11, ASC continued delivering science-based simulation tools for annual assessments and next-generation LEPs, focusing on improved physics, fidelity, and calculations in support of DSW and the National Code Strategy. The methodology for predictive capability assessment was demonstrated in FY11 for a limited set of simulations common to both physics laboratories. ASC assessed the ability to simulate full system performance near thresholds where data are sparse. Cielo, the capability HPC

[3] *2010 Nuclear Posture Review Report*, April 2010, p. 42.

[4] *FY11 ASC Program Plan* for past, present, and planned contributions to the SSP.

system, replaced Purple and began operation. Cielo was upgraded to 1.37 petaFLOP/s, the system was accepted, and the second round of the Capability Computing Campaign (CCC) was initiated. Sequoia, the advanced architecture system, received a "go" from the program's "go/no go" review. Contract negotiations and delivery schedule have been finalized. As of September 30, 2011, facility site preparation work was well underway, and design and procurement of the parallel file system was complete. Building on the FY10 Scalable Applications Project (SAP) milestone, the SAP effort has extended the knowledge base, documentation, and training to enable tri-lab code teams to use Sequoia. Successful procurement of the next-generation tri-lab Linux capacity clusters (TLCC2) and the associated common user environment milestones were executed across the three NNSA laboratories. Initial investments were made towards the proposed joint DOE Office of Science (SC) and NNSA exascale initiative. Proposal responses to ASC's Request for Information were evaluated in preparation for the next round of Predictive Science Academic Alliance Program (PSAAP) engagements.

In FY12, ASC released a Request for Proposal and reviewed proposals submitted for the PSAAP II program. Cielo, the new capability machine in operations at LANL, entered the General Availability phase in September 2012. All 96 racks of Sequoia were delivered in May 2012 and were integrated on the unclassified network by the end of the fourth quarter. Integration and operation of TLCC2 systems deployed across the tri labs were implemented by end of the fiscal year. FastForward investments, co-funded with DOE Office of Science Office of Advanced Scientific Computing Research to accelerate extreme-scale computing R&D technologies, were awarded to AMD, IBM, Intel, Nvidia, and WhamCloud. On the Physics and Engineering Models (PEM) front, a Level 1 milestone to advance capabilities for annual assessments and resolution of SFIs associated with early-phase primary implementation progressed as planned. Verification and validation assessment of improvements in primary performance codes for boost continued. The V&V subprogram also worked on the strategy of common modeling to validate improvements in support of the National Boost Initiative. The Integrated Codes (IC) program element concentrated on preparations for the next predictive capability framework (PCF) peg post and on a Level 1 milestone for simulating late-time primary implosion and initial explosion, as well as on scalability enhancements targeting future computing platforms.

In FY13 and beyond, ASC will be focusing on strengthening the science basis and driving down uncertainties for weapons simulations to a degree that NNSA can ultimately, and credibly, claim predictive capability; instituting a robust, formalized peer review system; increasing the number of production computing cycles to support increased use of simulation in stockpile activities and reliance on uncertainty quantification (UQ) in weapons decisions; and pursuing next-generation computing to meet time-urgent, future predictive science capability needs as documented in the *ASC Platform Strategy* and the *Predictive Capability Framework*.

III. Accomplishments for FY11–FY12

ASC accomplishments from quarter 4, fiscal year 2011, and through quarter 3, fiscal year 2012, are reflected below for the Computational Systems and Software Environment (CSSE) and Facility Operations and User Support (FOUS) sub-programs.

Computational Systems and Software Environment

LLNL Accomplishments for Computational Systems and Software Environment

- Delivered, sited, and performed initial integration of the Sequoia system, which appeared at the top of the Top500 list and the Green500 list, and tied for first on the Graph500 list
- Integrated the TLCC2 systems
- Integrated and tested pre-production products for Sequoia, including the development environment, Simple Linux Utility for Resource Management (SLURM) port, and ZFS-based Lustre file system
- Released enhancements to the Lorenz "MyLC" HPC Web portal, providing users with a powerful, personalized, application portal to improve access to Livermore Computing (LC) resources
- Contributed to post-petascale planning, including leading the FastForward procurement process, interacting with co-design centers and vendors, and collaborating with IC teams to evaluate next-generation technologies

LANL Accomplishments for Computational Systems and Software Environment

- Installed Luna, the LANL TLCC2 capacity system, and integrated the system into the classified computing environment
- Fully deployed the LANL production visualization cluster, Viewmaster2; currently being used by LANL scientists
- Released and deployed a production version of Parallel Log File System (PLFS) that has passed extensive regression testing and that is functionally correct when used with a major production simulation code
- Restructured projects to support co-design; initial proxy applications have been developed to foster interactions between IC developers, computer scientists, and computer vendors

SNL Accomplishments for Computational Systems and Software Environment

- Developed and applied a methodology for comparing mini-application (mini-app) performance to the performance of the complete application it is meant to represent
- Provided state-of-the-art test beds and analyzed bottlenecks of key mini-apps on these test beds; made recommendations to algorithm developers and computer architects based on this work
- Further expanded the capabilities of the Structural Simulation Toolkit (SST); integrated a disk model into SST/macro for check pointing, file input/output (I/O), and in situ analysis work and also enhanced the user interface; added scheduler, resiliency, and SST/macro models into SST/micro
- Integrated the ParaView co-processing library into SIERRA, giving SIERRA analysts and developers a flexible option for analyzing and visualizing data during a running simulation
- Demonstrated a characterization and feedback system for dynamic resource-aware computing by continuously rebalancing the workload of a SIERRA job while running on 10,112 processing elements on Cielo

Tri-Lab Accomplishments for Computational Systems and Software Environment

- Completed Cielo production capability readiness milestone, ensuring that machine accessibility and integration and operational support are in place and completed with required testing of user applications, system reliability, and application and I/O performance
- Created tri-lab working groups to identify and study issues related to next-generation computing technology on the path of exascale

Facility Operations and User Support

LLNL Accomplishments for Facility Operations and User Support

- Deployed the HPC Enclave, firewalls, and virtual private network configurations
- Deployed the Sequoia facilities (tertiary cooling loop commissioned, 10 MW of power distributed to 96 racks, data tray, rack stand re-design and deployment)
- Deployed the Sequoia network infrastructure (platform and storage area network (SAN))
- Integrated the TLCC2 systems (Zin, Cab, Rzmerl) and the network, including power re-work and redistributing BlueGene/L power to Zin
- Upgraded CHAOS 5 to all capacity systems in the center

LANL Accomplishments for Facility Operations and User Support

- Completed L2 milestone 4473, *Production Readiness for the Luna Capacity Computing System*; the Luna cluster adds over 500 teraFLOP/s to the LANL secure computing network
- Completed site preparation, analysis, and testing for Sanitary Effluent Reclamation Facility (SERF) water integration at the Strategic Computing Complex (SCC); SERF recycles 120 million gallons a year of sanitary effluent for industrial reuse and is a major contributor for DOE/NNSA meeting requirements for recycling and reducing water use
- Completed 60 percent of the design phase for the SCC upgrade in support of future ASC water-cooled platforms
- Deployed the PLFS on all ASC platforms; PLFS increases dump time performance for n–1 workloads by converting an n–1 request into n–n

SNL Accomplishments for Facility Operations and User Support

- Continued New Mexico Alliance for Computing at Extreme Scale (ACES) user support collaboration for Cielo and further integrated ticket management and Tier 2 and Tier 3 support
- Installed over 1 petaFLOP/s aggregate computing resources to serve the nuclear weapons and National Security Computing Center (NSCC) customer base; each of three systems were obtained through the TLCC2 procurement and are using the common Tripod Operating System Software (TOSS) software environment
- Made the DisCom network fully redundant (no single fiber paths existed) and upgraded the network with Cisco Routers and 10 Ge encryptors
- Adopted the architecture for data transfers in use at LANL and LLNL; SNL Lynx platforms mount global file systems, have access to High-Performance Storage System (HPSS) systems, and act as the transfer point for data movement between sites across the DisCom network
- Continued operations support of the TLCC1 systems and hierarchical archive platforms on restricted and classified networks, upgraded global file systems, and began initial activities on the Intersite HPC network interconnecting the three labs

Academic Alliances

- Caltech completed 75 percent of the annual UQ campaign, bringing together the Eulerian and Lagrangian solvers, the extended material models, the Legacy-OUQ protocol, and the hypervelocity legacy SPHIR data.
- Purdue University made a successful first-pass at Bayes network analysis of pull-in, pull-out, dielectric charging, and creep in MEMS employing heterogeneous data on

frog-leg, cantilever, and Center for Prediction of Reliability, Integrity, and Survivability of Microsystems (PRISM) devices.

- Stanford University demonstrated levels of performance equal to those of hand-tuned code for four different applications written in the Liszt domain-specific language, while retaining automatic portability to distributed- and shared-memory machines, as well as GPUs.
- University of Michigan enabled the CRASH laser package to trace rays in 3D and deposit energy in 2D or 3D.
- University of Texas integrated FIN-S with updated physics models in turbulence, transport, and chemistry as guided by validation.

IV. Product Descriptions by the National Work Breakdown Structure

Computational Systems and Software Environment (WBS 1.5.4)

Commodity Technology Systems (WBS 1.5.4.8)

Production Planning and Integration (LLNL)

Accomplishments in FY12

- Received, integrated, and provided production support for new TLCC 2 systems
- Provided production support for TLCC1 systems

Planned Activities in FY13:

- Monitor computer industry developments for opportunities to enhance capacity computing and the associated infrastructure
- Provide production support for TLCC2 systems

Computing Platform Integration and Deployment (LANL)

Accomplishments in FY12

- Installed Luna, the TLCC2 capacity system, and integrated the system into the LANL classified computing environment
- Installed Moonlight, a hybrid TLCC2 capacity cluster (general-purpose computing on graphics processing units (GPGPU)-enhanced TLCC2 cluster), into the LANL unclassified computing environment
- Completed the ASC L2 review for Luna production readiness

Planned Activities in FY13:

- Upgrade the Luna system with two additional scalable units (SUs)
- Continue to operate Luna and the other capacity systems in both the classified and unclassified computing environments
- Complete the integration of Moonlight, a TLCC2 system with GPGPUS, into the LANL unclassified computing environment; work with the applications community for identification of additional software tools for the hybrid architecture
- Develop requirements for the next round of Commodity Technology Systems with tri-lab design team

Advanced Technology Systems (WBS 1.5.4.3)

Sequoia Tri-Lab Advanced Technology Platform (LLNL)

Accomplishments in FY12

- Delivered and installed all 96 racks of Sequoia
- Issued system accounts to early users, who have run jobs
- Achieved 16.32 petaFLOP/s on the Linpack benchmark using 1,572,864 cores, coming out on top of the Top500 list and tied with ANL's Mira system (both BlueGene/Q technology) for the top slot on the Graph500 list (measures performance of data-intensive supercomputer applications)
- Earned #1 on the Green500 list, which ranks the most energy efficient supercomputer systems

Planned Activities in FY13:

- Make Sequoia available for unclassified science runs for three to four months
- Transition to the classified network and complete security testing
- Place into production on the classified network for tri-lab classified work

BlueGene/P and BlueGene/Q Research and Development (LLNL)

Accomplishments in FY12

- Collaborated with IBM to develop BlueGene/Q human heart simulation
- Guided design of Lightweight OpenMP runtime
- Collaborated with IBM to develop power measurement application programming interface (API) for BlueGene/Q
- Explored parallelization of CPLEX dynamic programming framework

Planned Activities in FY13:

- Complete BlueGene/Q research and development (R&D) contract
- Develop, deploy, and test complete Lightweight OpenMP prototype
- Identify and evaluate techniques to reduce BlueGene/Q power consumption
- Refine and optimize human heart simulation

Hyperion Test Bed (LLNL)

Accomplishments in FY12

- Deployed new high performance disk storage systems into the Hyperion environment
- Procured a technology refresh of Hyperion phase 1 servers, including high performance storage class memory
- Supported the scalability testing of the Lustre file system, high performance storage hardware, system software, and application middleware

Planned Activities in FY13:

- Deploy technology refresh of the Hyperion phase 1 servers
- Procure and deploy a technology refresh of the Hyperion phase 2 servers, including high performance storage class memory, InfiniBand cluster interconnect, and data intensive capabilities
- Explore new software models to evaluate the use of high performance storage class memory and the design impacts of storage class memory on future system software and hardware architectures
- Continue to support scalability testing on system software, middleware, storage, and file systems

Alliance for Computing at Extreme Scale Trinity Advanced Technology System (SNL, LANL)

Accomplishments in FY12

- Created the Trinity Advanced Technology System Project
- Completed the vendor market survey for Trinity
- Initiated planning for a joint procurement process for Trinity and NERSC-8 through ACES and Lawrence Berkeley/NERSC
- Completed Trinity CD0

Planned Activities in FY13:

- Complete Trinity CD1
- Release Trinity/NERSC-8 RFP
- Evaluate Trinity/NERSC-8 RFP responses and make selection

Alliance for Computing at Extreme Scale Cielo Capability Computing Platform (SNL, LANL)

Accomplishments in FY12

- Completed the Cielo system integration
- Completed the ASC Level 2 milestone, *Cielo Production Capability Readiness*
- Completed the CCC2 on Cielo
- Completed a risk mitigation plan for the Cielo file system; transitioned to using the Lustre file system on Cielo

Planned Activities in FY13:

- Complete the Cielo CCC3 projects and start the CCC4 projects
- Complete the ASC Level 2 milestone for Lustre production on the Cielo system
- Continue to run Cielo in production capability mode

Next-Generation Computing Technologies (WBS 1.5.4.9)

Next-Generation Computing Enablement (LLNL)

Accomplishments in FY12

- Participated in exascale planning activities, representing NNSA/LLNL at joint ASC/ Advanced Scientific Computing Research (ASCR) meetings, Exascale Centers of Excellence, and various workshops
- Performed exascale vendor technology evaluations, including reviewing exascale request for information responses
- Interacted with co-design centers and ASC IC teams to discuss and evaluate next-generation technologies (many integrated core (MIC) processor, programming models, graphics processing units (GPUs), correctness tools, and power tools)
- Facilitated the FastForward procurement award process, leading to five contracts

Planned Activities in FY13:

- Continue to participation in next-generation computing planning activities, including the development of a tri-lab plan for a next-generation software environment
- Conduct co-design activities with ASC and ASCR co-design centers and vendors, and research and evaluate next generation technologies
- Explore post-petascale topics, including novel threading environments and performance characterization and other code improvement opportunities for next-generation systems
- Provide technical coordination for the FastForward contracts
- Participate in requirements gathering for next-generation environments, including storage, data analysis and visualization, development environment, and data workflow and data management

Systems Requirements and Planning (LANL)

Accomplishments in FY12

- Participated in the ASC platform strategy and the development of the new platform funding model; this new model transitions the ASC platform strategy from three classes of systems (capability, capacity, and advanced architectures) to two classes of systems (advanced technology systems and commodity technology systems)

Planned Activities in FY13:

- Plan for the Trinity platform as part of the NNSA ASC Platform program
- Identify facility, power, memory, and file system requirements for exascale systems
- Participate in site-wide planning for power upgrades for future systems
- Plan infrastructure to support pre-exascale and exascale systems

Co-Design through Mini-Applications (LANL)

Accomplishments in FY12

- Released a two-dimensional (2D) Shock Hydro CUDA/graphics processing unit (GPU) mini-app (and document) as open source
- Extracted a mesh traversal mini-app from CHICOMA Lagrangian hydrodynamics code
- Developed an experimental mesh traversal mini-app in high-level programming language
- Designed and built a new standard thread-capable memory request and granting scheme that allows for a set of accelerator cores to exchange data with the central processing unit (CPU) in a fine-grained, low-latency manner
- Designed and built solid-state drive (SSD)-based hardware check pointing system that provides check pointing to an application as an independent, local, high-speed hardware service

Planned Activities in FY13:

- Work with ASC code teams to select primary areas for mini-app development, and assist code teams in developing specification documents and reference codes
- Study mini-apps and their effective implementation on next-generation architectures; work with ASC code teams to identify and implement useful abstractions for multi-physics applications on extreme-scale computing systems
- Continue work on refactorization of ASC codes, moving new code capabilities into production

Programming Models for the Next-Generation Scientific Computing Environment (LANL)

Accomplishments in FY12

- Designed, developed, and tested an architecture agnostics, compiler-based profiling capability capable of tracking application memory and performance characteristics related to changing/future system architectures
- Studied the impact of NUMA issues on current codes and the implications of underlying architecture details
- Studied the implementation and characteristics of stencil-based calculations with a goal of quantifying performance and identifying bottlenecks (for example, balance of code readability and maintainability versus performance)

Planned Activities in FY13:

- Collaborate with the co-design through mini-apps and visualization and data analysis projects to investigate the impact of using existing APIs to program next-generation architectures for both physics and data-intensive applications
- Develop techniques and tools for the analysis and potential transformation of code structures given the dynamic behavior of select applications; the goal is to assist in the identification of target code elements for acceleration on new architectures, and explore how these findings can be leveraged to define new programming models
- Develop techniques and tools for the analysis of data movement through the memory hierarchy and its impact on application performance and scaling (including the I/O subsystems); explore how these findings can be leveraged to define new programming models
- Participate in and support the activities of the ASC CSSE exascale planning working group on programming models

Co-Design Enablement with Computer Science (SNL)

This is a new project for FY13.

Planned Activities in FY13:

- Create use cases to assess and specify development and deployment approaches for critical components in ASC computing environment
- Evaluate and quantify potential technical gaps or issues and where such gaps exist, define a prioritized approach to closing them
- Create baseline set of candidate proxy architecture definitions

Advanced Systems Technology Research and Development (SNL)

Accomplishments in FY12

- At Supercomputing 2011, successfully demonstrated simulation results from SNL's 42-node cluster with new Intel MIC processors
- Demonstrated execution of the FE proxy mini-app using message passing interface (MPI) and OpenMP on SNL's full cluster of 84 Knights Ferry cards
- Ported OpenCL version of mini-apps to next-generation AMD/Fusion Trinity Accelerated Processing Unit (APU) achieving a two-fold performance improvement
- Completed Phase II upgrade of AMD/Fusion cluster to Trinity APUs and custom integrated power measurement boards
- Completed Phase III upgrade of Intel/MIC cluster to Knights Corner prototype cards

Planned Activities in FY13:

- Generate initial validation data for comparison with SST architectural simulations driven by proxy applications
- Continue exploration of performance capabilities of advanced architectures with focus on vectorization, threading, and internode strategies and power consumption
- Provide feedback to industry collaborators on the impact of specific hardware capabilities for future systems
- Integrate power measuring capability on commodity hardware

Application Performance Analysis for Next-Generation Systems (SNL)

Accomplishments in FY12

- Provided national leadership in defining the role of mini-apps for exascale through participation in conferences, workshops, and symposiums
- Collaborated on co-design with Intel's MIC Architecture Team at the Intel EPOCH workshop
- Completed Level 2 milestone, *Characterize the Role of the Mini-Application in Predicting Key Performance Characteristics of Real Applications*

Planned Activities in FY13:

- Work with co-design centers and exascale activates in applying mini-apps and proxy applications
- Develop and apply methods to use mini-apps and proxy applications and define their predictive capabilities
- Integrate in the use of SST for exploring the predictive capabilities of mini-apps on future platforms
- Provide support to the Trinity design team in defining the role of mini-apps in the procurement evaluation and acceptance phases of the acquisition

Heterogeneous Computing (SNL)

Accomplishments in FY12

- Developed the Kokkos-Array performance-portable programming model, which integrates many-core data parallel algorithms
- Demonstrated many-core performance-portability on NUMA and GPU devices with linear algebra and finite element unit-test computational kernels
- Developed hybrid parallel nonlinear thermal conduction and explicit solid mechanics finite element proxy-applications to demonstrated MPI + Kokkos-Array programming model
- Evaluated performance of hybrid parallel proxy applications on SNL's Cray XK6 test bed using both NVIDIA GPU and NUMA CPUs for fine-grained parallelism
- Characterized performance concerns associated with fine-grained parallelism on NUMA CPUs

Planned Activities in FY13:

- Develop Kokkos-Array runtime "back end" for Intel MIC many-core/accelerator
- Evaluate performance unit-test kernels and proxy applications on Intel MIC-based test bed
- Research programming model characteristics required to support instruction-level vectorization within fine-grained parallel algorithms
- Research performance interactions between MPI and Kokkos-Array levels of parallelism
- Research MPI + CPU + accelerator heterogeneous parallelism
- Research hierarchical, heterogeneous domain decomposition to address NUMA performance concerns and work decomposition in a MPI + CPU + accelerator strategy

System Software and Tools (WBS 1.5.4.4)

System Software Environment for Scalable Systems (LLNL)

Accomplishments in FY12

- Deployed TOSS 2.0, including Red Hat Enterprise Linux (RHEL) v6.2-based, TLCC2 hardware support, Lustre v2.1, SLURM v2.3, and Moab v6.1
- Deployed Infiniband SAN support for Sequoia
- Ported SLURM to Sequoia in integration/test

Planned Activities in FY13:

- Continue ongoing TOSS software development and support, including tracking RHEL 6 releases and integrating into TOSS releases as appropriate
- Continue ongoing development and support of Moab and SLURM, including tracking releases of Moab and SLURM and integrating into TOSS releases as appropriate
- Implement new system software development activities for efficient operation of ~10,000 node scale to address increasing commodity cluster sizes, increasing core counts, and the ever-present I/O bottleneck, including:
    - Investigation of a flexible communication network to serve multiple uses, such as integration of node/system monitoring, support for debugging/profiling tools, scalable MPI job launch, and support for security, scalable authentication, encryption, and privacy/integrity
    - Investigation of true generic and heterogeneous resource scheduling (for example, nodes, sockets, cores, threads, memory, GPUs, disk, and I/O bandwidth)
    - Investigation of alternate network technologies for clusters and SAN

Applications Development Environment and Performance Team (LLNL)

Accomplishments in FY12

- Deployed and tested the early BlueGene/Q development environment on Sequoia
- Demonstrated Stack Trace Analysis Tool (STAT) scalability to 768,000 MPI processes on 48 racks of Sequoia
- Co-designed the Sequoia code development tools infrastructure and successfully demonstrated it in both vendor and Applications Development Environment and Performance Team (ADEPT) code correctness tools
- Deployed the development environment for TLCC2 and TOSS 2.0
- Supported migration of applications teams to production use of the TLCC2 systems

Planned Activities in FY13:

- Continue code development environment support on all LLNL ASC platforms
- Assist tri-lab integrated design code teams on Sequoia as part of SAP activities
- Develop tool infrastructures to improve both scalability and performance of applications and the code development environment
- Enhance capabilities for the code correctness tool suite on TLCC2 and Sequoia
- Provide dedicated performance tuning expertise for key applications on Sequoia

High Performance Computing Systems Research (LANL)

Accomplishments in FY12

- Developed system for optimizing latency and throughput for spawning processes on massively multicore processors
- Developed techniques leading to near-noiseless Linux environments
- Organized Fault Tolerance for HPC at Extreme Scale Workshop in association with the 42nd annual IEEE/IFIP International Conference on Dependable Systems and Networks
- Developed FILESIM, a parallel file system discrete event simulator, and validated it against a TLCC2 test bed cluster representing a Panasas parallel file system

Planned Activities in FY13:

- Continue developing redfish library to support exascale services; focus on services that can address issues for upcoming Trinity timeframe and scale, that is, power management; reliability, availability, scalability (RAS); and distributed job launch
- Investigate and develop a Parallel file system layer busting as a simple alternative to the POSIX interface that is built on top of the low-level, highly parallel services already provided by the Ceph parallel file system, to directly implement HDF5 scientific data format
- Mature the prototype soft error fault injector into a pre-production research tool for resilience; experiment on fault response to simulated soft errors in applications of interest to the ASC program
- Investigate application performance on multi and many-core nodes; investigate a single node for how to optimize NUMA layout within system and application memory bandwidth limits; identify the memory bandwidth requirements of applications via instruction stream analysis and measure with PAPI and Pintool; verify and validate on multicore and many-core computers
- Create version simulate Lustre after validating FILESIM with Panasas; use this to design comparisons for large-scale file systems

Application Readiness (LANL)

Accomplishments in FY12

- Performed Luna, Mustang, and Moonlight acceptance testing, including techniques for screening node performance variability; integrated LLNL and SNL applications to Gazebo test suite
- Analyzed and developed workarounds for a variety of high-impact application problems, including damaged checkpoints, I/O challenges on several integrated scientific codes, and performance variability problems on Cielo
- Developed prototype job log analysis for improved diagnosis of application-system problems

Planned Activities in FY13:

- Support users of Luna, Mustang, Moonlight, and other recently integrated systems by tackling hard-to-diagnose problems, typically involving the interaction of applications with multiple aspects of the computational environment
- Work with other projects including Co-Design through Mini-Applications, Programming Models, and Code Strategies for Emerging Platforms to assist code development teams with migration towards the use of hybrid programming to exploit the GPU-based elements of a new TLCC2 system
- Provide support to production applications on capacity platforms
- Assist system management personnel with problem investigation and resolution
- Develop application observation techniques for automated scanning of key application progress and success/failure indications

Software Support (LANL)

Accomplishments in FY12

- Leveraged the work of the Common Computing Environment (CCE) project to develop plans toward exascale tools solutions
- Supported and deployed the initial deliverable from the scalable debugger work with LLNL/Rogue Wave (TotalView)
- Completed the Open MPI port to the Cray XE6 platform
- Continued local support for tools and scheduling tool workshops between users and tool providers

Planned Activities in FY13:

- Continue building an MPI support capability by engaging the community support model; focus will be on Open MPI development targeted to tri-lab needs
- Continue performance analysis support capability; focus will be on Open|SpeedShop (O|SS)/component-based tool framework (CBTF) integration in application analysis and new analysis needs
- Increase memory analysis and MPI performance/correctness tools and support
- Continue building a stronger debugger support capability; work with TotalView on tool deployments and enhancements; work with LLNL debugger capability through STAT and subset debugging
- Continue involvement with DOE NNSA and DOE/SC laboratories in exascale tool planning and development programs
- Continue development support of the CBTF, a common, scalable framework to ease implementation, integration, and deployment of software development tools

Software and Tools for Scalability and Performance (SNL)

Accomplishments in FY12

- Released joint version of Kitten V1.3 and Palacios V1.3
- Developed tool to measure the memory bandwidth of a running application with correctness validated; this tool will be used to quantify the benefit of continuous data placement optimizations
- Ported Palacios lightweight virtual machine monitor to Cray XK6 compute node Linux operating system
- Delivered early prototype of commodity power measurement board as part of AMD Fusion Trinity early release test bed node
- Demonstrated on-chip power monitoring capability on prototype Trinity APU

Planned Activities in FY13:

- Integrate energy measurement capability on commodity hardware platform
- Develop initial power API
- Integrate dynamic shared library support into Cray XT job launch mechanism
- Test and optimize large-scale dynamic shared library usage
- Support high-performance networking communication between virtual machines
- Provide support for lightweight kernel measurement capabilities for application runtime characteristics of energy usage, memory bandwidth, and CPU utilization

Resilience (SNL)

Accomplishments in FY12

- Deployed Splunk-HPC log analysis servers for TLCC, Red Sky, TLCC2, and Cielo
- Tested on Cielo and other systems the OVIS sensor-based monitoring
- Implemented fault simulation framework implemented in SST (Macro)

Planned Activities in FY13:

- Develop persistent data store object, including integrated vertical demonstration of prototype application recovery from local fault via local recovery using persistent storage
- Provide requirements and perform analysis of basic features needed to support local recovery from persistent storage, classified by system layers
- Enhance dynamic scheduling and planning for task-based execution model
- Develop resilient PDE solvers, redesigning algorithms to make progress despite faults
- Develop principles by which parallel algorithms for solving PDEs can be modified to make them robust to silent (as well as detected) faults using algorithmic properties
- Understand faults and their effects using Splunk for HPC
- Enhance functionality of the Splunk log analysis toolkit in three areas: 1) additional support for analysis of Lustre file system logs; 2) develop models with UIUC on Cray XE6 associating fault types, occurrence rates, propagations, and their effects on jobs; and 3) implement the "nodeinfo" anomaly detection algorithm within Splunk

System Simulation and Computer Science (SNL)

Accomplishments in FY12

- Added fault injection and power modeling capability to SST/macro
- Improved the network congestion model to achieve more scalable performance and model switch buffer effects
- Supported SST/micro, OMNeT, and SystemC back ends with SST/macro
- Added real C and Fortran MPI interfaces to facilitate skeletonizing and running applications within the simulator
- Grew the suite of application skeletons that ship with SST/macro to include network-attached storage (NAS)-lu, sweep3D, GTC, and miniGhost

Planned Activities in FY13:

- Demonstrate hierarchical validation of SST/macro against real machine, using UQ for determining accuracy in scaling beyond current capabilities
- Support fault resilience work by developing and exploring implementations of next-generation middleware, runtime, and libraries
- Support and assist in the development of experiments for co-design centers
- Document lessons learned from FY12 release of SST that integrated micro and macro components running on four late-model operating systems and compilers, and apply them to this year's feature release
- Develop validation framework for SST/micro and its components

Scalable, Fault-Resilient Programming Models (SNL)

Accomplishments in FY12

- Developed 1D cellular automaton to explore bounding behavior of task-DAG programming model approach in presence of faults and analyzed scalability and resilience behavior in SST and on Cielo
- Developed a task-DAG API for use with miniFE to explore research challenges for a realistic problem
- Developed "robust stencil" for simple hyperbolic PDE, and generated computational experimental results showing potential for resilience to silent errors

Planned Activities in FY13:

- Formalize the methodology for comparing performance of task-DAG programming models to SPMD-style applications; compare at (simulated) scale and in the face of failures
- Continue exploration of the programming model design space to include 1) applications more difficult to load balance, 2) new failure response strategies, and 3) models of failure histories
- Ensure the ability of the programming model to provide the expected improvements in fault recovery while maintaining or improving load balance by examining realistic workloads
- Create initial metrics for development effort and performance profile analysis for non-MPI programming models
- Validate simulator performance predictions with realistic application workloads using Cielo and/or other current ASC platforms; compare mini-app performance under different fault recovery strategies

Input/Output, Storage Systems, and Networking (WBS 1.5.4.5)

Archive Storage (LLNL)

Accomplishments in FY12

- Supported the Level 2 milestone, *HPC Enclave*, by instantiating archive security zones within a single namespace
- Deployed HPSS 7.3.3.patch2 and new Linux mover clusters and cache disk
- Evaluated, selected, and procured new Linux-based HPSS Core Server platforms and metadata disk subsystems
- Tested HPSS 7.4.1 with Redundant Array of Independent Tapes (RAIT)
- Began development of HPSS 7.P (a.k.a. Panda), featuring an architecture utilizing partitioned metadata

Planned Activities in FY13:

- Continue ongoing HPSS software development and support, including 1) provide ongoing maintenance and customer support for HPSS 7.3; 2) finalize testing of and release HPSS 7.4.1 with RAIT; and 3) continue development of HPSS 7.P (a.k.a. Panda), featuring partitioned metadata
- Deploy Linux-based HPSS core server platforms and metadata disk subsystems
- Deploy Solaris-based platforms and Oracle's Automated Cartridge System Library Software (ACSLS) 8.0 for Oracle-based tape libraries
- Support the tri-lab ASC *Archive Growth Sustainability Study*
- Provide ongoing support of currently deployed archival storage systems, including selection, deployment, support, and maintenance of all archival storage hardware and media, customer and interface support, ongoing tech refresh, and data stewardship

Parallel and Network File Systems (LLNL)

Accomplishments in FY12

- Integrated and deployed pre-production ZFS-based Lustre file system for Sequoia
- Participated in development and engineering effort to provide a canonical file system layer within Lustre, allowing alternative file system underpinnings
- Provided Lustre file system performance and scalability enhancements in support of Lustre 2.1; deployed release on eight production file systems
- Changed NAS architecture for LLNL HPC enclave
- Maintained and supported Lustre, NAS, and global parallel file system (GPFS) parallel file systems, including middleware and higher level I/O libraries for users

Planned Activities in FY13:

- Field production-ready Sequoia file system in classified environment
- Complete development of porting layer with Lustre, allowing multiple file system back-ends
- Deploy a ZFS-enabled release of Lustre into TOSS
- Participate in deployment and development of the next Lustre feature release

Networking and Test Beds (LLNL)

Accomplishments in FY12

- Installed a 40GE test host that connects to the 100GE wide area network (WAN) test circuit; the host includes the latest technology such that bandwidth test could potentially approach wire speed rather than the previous 40 percent of capacity
- Evaluated and integrated low latency 40GE Ethernet switch for data center Ethernet
- Evaluated several NAS accelerators and negative caches
- Updated support for ConnectX3 fourteen data rate-based cards in Open Fabrics Alliance
- Evaluated open source software for managing and performance tuning solid state disks

Planned Activities in FY13:

- Manage Computational accelerator and test the Nvidia Kepler and Intel Xeon Phi
- Evaluate alternative Lustre hardware platforms with the goal of reducing total cost of ownership
- Evaluate cluster management tools at scale
- Network block device reliability and performance testing at scale
- Continue study of developing load balancing and multipath routing in support of congestion avoidance for previously listed networks
- Apply testing results to optimize the functionality, performance, reliability, manageability, and security of the I/O services supporting these computing systems
- Evaluate JBOD for use with Lustre
- Test ZFS ZIL and L2 buffer with Lustre
- Evaluate open source Lustre management tool
- Evaluate Quad socket EP-based Intel nodes for large-scale memory applications and testing of HPSS core services

Archival and File Systems (LANL)

Accomplishments in FY12

- Released a production version of PLFS that has passed extensive regression testing by LANL and EMC (LANL's CRADA partner), and that is functionally correct when used with a major production simulation code
- Released metadata and static file tree information for LANL file systems to be used by researchers and others in evaluating file system and archive efforts
- Published papers on PLFS, storage challenges, and I/O performance for IEEE, MSST, and internal use
- Released promised version of parallel storage interface (PSI) and installed promised versions of HPSS to LANL production systems
- Added host-to-host and local transfer features to beta PSI

Planned Activities in FY13:

- Co-develop PLFS with EMC under our CRADA with a focus on run-time efficiency and scalability for currently installed systems
- Verify parallel file system software releases (Panasas and Lustre) and assist with file system stand-up of incoming machines
- Support the Applications Readiness project with any I/O issues and general production needs
- Support the production archive by assisting users and altering software to address use issues

Archival Storage (SNL)

Accomplishments in FY12

- Released 7.3.3 of HPSS with new features for user-defined attributes and checksum APIs, as well as a number of improvements for scalability
- Moved DISCOM parallel file transfer protocol (PFTP) into supported tree

Planned Activities in FY13:

- Add RAIT support
- Identify and remedy bottlenecks in existing code
- Design and prototype v7.p with partitioned database and multiple nodes on a single system; target is 5000 file creates per second

Scalable Input/Output Research (SNL)

Accomplishments in FY12

- Released first open-source version of TRILINOS I/O Support (Trios) software, which includes CSSE data services, netCDF staging, and SEACAS
- Performed SICAIDA work with a project nominated for an NNSA Defense Programs Award of Excellence
- Demonstrated early prototype of SIROCCO file system at SC11
- Productized Gibraltar RAID (GPGPU RAID)

Planned Activities in FY13:

- Evaluate and compare in-situ and in-transit analysis of CTH data (part of an FY13 Level 2 milestone)
- Identify and begin development of data services for SIERRA codes
- Develop and demonstrate an expandable "burst buffer" that adds compute-node resources as needed to support bursty I/O writes
- Demonstrate an early prototype of SIROCCO that includes storage servers that use gossip protocols to exchange information and log-based on-disk storage
- Demonstrate data services for GPU accelerators

Scalable Interconnects for Extreme-Scale Tightly Coupled Systems (SNL)

Accomplishments in FY12

- Completed an initial analysis of the sensitivity of applications to network injection bandwidth
- Released an initial reference implementation of Portals 4.0 for InfiniBand networks
- Completed a study of Portals 4.0 triggered operations for MPI non-blocking collective operations

Planned Activities in FY13:

- Release a unified reference implementation of Portals 4.0 for shared memory and InfiniBand networks that provides asynchronous progress suitable for many-core processors
- Complete an analysis of application sensitivity to network injection bandwidth and link bandwidth on more than four thousand cores on Cielo for several applications
- Develop, test, and deploy an enhanced version of OpenMPI for Portals 4.0 that contains optimized implementations of blocking and non-blocking collective operations using triggered operations

Post-Processing Environments (WBS 1.5.4.6)

Scientific Visualization (LLNL)

Accomplishments in FY12

- Deployed new one-node PowerWall driver to classified Visualization Work Center, replacing a multi-node cluster and using personal computer over internet protocol technology
- Deployed and supported suite of visualization and data analysis tools on LLNL clusters, including GPU-specific software on the 216-node Edge data analysis cluster
- Provided operational support for all visualization facilities, including multiple PowerWalls; provided consulting and development support for data analysis as well as support for presentations
- Performed R&D in visualization and data analysis, mentored students, and participated in post-petascale planning
- Designed and implemented an initial visualization system for Lorenz using the new Web graphics library standard that enables native rendering capabilities through the Web browser

Planned Activities in FY13:

- Continue to maintain the data analysis and visualization hardware platforms and software environment
- Provide operational support for all visualization facilities, including supporting projection equipment and facilitating the use of the data analysis clusters and associated storage
- Support large-scale data analysis and visualization activities, including supporting ASC scientists with the creation of visuals and movies for presenting and analyzing scientific data
- Exploit research results in data analysis and visualization to prepare for post-petascale analysis and to extend capabilities, including data compression and topological methods
- Expand WebGL visualization system for Lorenz by incorporating streaming and multi-resolution capabilities as well as advanced volume visualization techniques

Scientific Workflow and Data Management (LLNL)

Accomplishments in FY12

- Released versions 2.8 and 2.9 of Hopper and Chopper, featuring new options for efficiently dealing with massive datasets and directories, extensive usability improvements, and improved cross-lab transfer capabilities
- Released extensive enhancements to the Lorenz "MyLC" dashboard, providing users with a powerful and personalized HPC Web portal
- Released first production version of the Lorenz job management component, providing Web-based job submission, monitoring, and control.
- Released a beta version of a vertical application within the Lorenz application portal

Planned Activities in FY13:

- Release of new versions of Hopper and Chopper with architecture enhancements, enabling better support for multi-zone and multi-site installations, much reduced resource footprint for client-server operations, and numerous usability improvements
- Incorporate Web-based visualization into Lorenz, including as an application-specific examination mode in the job management component and as a general utility
- Extend Lorenz application portal to include one or more additional ASC-relevant applications
- Develop and release new MyLC dashboard portlets that provide users with additional insight into the computer center

Production Visualization (LANL)

Accomplishments in FY12

- Deployed LANL's production visualization cluster, Viewmaster 2
- Put into production the Los Alamos Weapons Conference Center, a conference room and museum of LANL devices equipped with state-of-the-art visualization tools, including touch screen capability
- Progressed on Trinity visualization planning
- Extended the capability to do high-performance ASC visualization to buildings outside of the SCC

Planned Activities in FY13:

- Finalize design and deploy new PowerWall Theater
- Continue to support and maintain production visualization systems
- Provide contract management and requirements specification, including facilities, visualization cluster, and EnSight contracts; support and maintain the EnSight software and help direct EnSight development activities under the new LANL EnSight development contract to Computational Engineering International
- Provide technical guidance on visualization needs for the forthcoming Trinity Advanced Technology System
- Work directly with designers in physics-based, iterative discovery process using the petascale visualization and data analysis enabled tool (EnSight) to promote new discoveries in weapons science in programs

Visualization and Data Analysis Research and Development Project (LANL)

Accomplishments in FY12

- Developed a portable, data-parallel visualization and analysis library (called PISTON); deployed an open-source release of this software, including efficient implementations of several fundamental visualization operators; and integrated it with visualization toolkit (VTK) and ParaView, including preliminary support for multi-node parallelism
- Created a framework and algorithms (unique values and first difference detection) for automatic feature detection for in situ analysis
- Demonstrated automatic in situ analysis in xRAGE, producing high temporal resolution imagery at run time for the 2D adaptive mesh refinement LCROSS simulation

Planned Activities in FY13:

- Explore, develop, and deploy the end-to-end integration of in-situ analysis and feature-extraction visualization on a variety of next-generation supercomputers to improve the understanding of massive scale results
- Develop a data-intensive post-processing visualization and analysis infrastructure that meets ASC programmatic needs by leveraging existing and new custom data-intensive software
- Participate in and support the activities of the ASC CSSE exascale planning working group on visualization and data analysis

Scalable Data Analysis (SNL)

Accomplishments in FY12

- Released flexible build for Catalyst, an in-situ analysis and visualization software library built on the VTK/ParaView core
- Integrated initial release of Catalyst into SIERRA code base, supporting basic use cases for in-situ analysis
- Conducted initial scaling testing of Catalyst integrated with a science code, running on 64k nodes of Cielo
- Completed two ParaView releases (versions 3.12.0 and 3.14.0) with improvements to the ServerManager and the Plugin framework; metrics show more than 100 active internal SNL users
- Delivered ParaView for production use on Cielo and supported CCC1 and CCC2

Planned Activities in FY13:

- Integrate first release of in-situ analysis capability (Catalyst) with SIERRA codebase, with production support
- Complete FY13 Level 2 Milestone, *Data Co-Processing for Extreme Scale Analysis* to develop and test multi-mode in-situ capabilities that reflect a real-world problem; team is testing at large scale on Cielo
- Deliver scalable analysis and visualization capabilities for Sequoia customers
- Participate and contribute to system design and user environment planning for Trinity
- Develop initial analysis and visualization algorithms on many core architectures; in conjunction with the scalable analysis community
- Continue *ParaView* releases, with production support in conjunction with Kitware, Inc.

Facility Operations and User Support (WBS 1.5.5)

User Support Services (WBS 1.5.5.2)

Hotlines and System Support (LLNL)

Accomplishments in FY12

- Provided ongoing support services for hotline operations, documentation, and training
- Analyzed two years of LC incidents and identified areas of improvement in hotline operations
- Deployed CRYPTOCard tokens to large subset of the LLNL user community in support of the HPC Enclave
- Developed Sequoia system documentation and training
- Deployed Confluence Wiki, new Staff Away, and System Support Matrix Web tools

Planned Activities in FY13:

- Use LLNL standards to deploy a redesigned Web infrastructure for LC
- Deploy new whiteboard, system status, and events tools for improved user communication
- Provide support to the new high performance computing innovation center (HPCIC) user community
- Assist users in the migration of applications to the BlueGene/Q hardware architecture
- Integrate Front Range and JIRA[5] for improved problem reporting and bug tracking
- Analyze and improve bank allocation procedures
- Continue to implement information technology infrastructure library (ITIL) best practices in hotline operations

[5] JIRA is a proprietary issue tracking product, developed by Atlassian, commonly used for bug tracking, issue tracking, and project management. The product name, JIRA, is not an acronym but rather a truncation of "Gojira", the Japanese name for Godzilla.

Integrated Computing Network Consulting, Training, Documentation, and External Computing Support (LANL)

Accomplishments in FY12

- Developed user support tools, capabilities, and infrastructure to support users in the ASC/LANL/ACES computing resources
- Performed data-supported analysis of key factors limiting end user compute availability
- Developed a framework prototype for continuous system software testing
- Developed several components for a Distributed Data Services (DDS) middleware layer and implemented them to process disparate data sources across capability and capacity platforms

Planned Activities in FY13:

- Develop new end-user interfaces to deliver comprehensive job information using DDS middleware
- Deploy ITIL-based modules to restructure ticketing system into problem and issue hierarchy
- Perform ongoing user support for users of ASC/LANL/ACES computing resources

User Support (SNL)

Accomplishments in FY12

- Provided coordinated and tiered user support for SNL's ASC resources and, in partnership with LANL, for ACES resources
- Developed integrated, "end-to-end" HPC user training content
- Improved active content for the SNL HPC and ACES user-focused Web sites
- Developed agreements to ensure maintenance and improvements of collaborative user support tools and teams.

Planned Activities in FY13:

- Provide user support for SNL and tri-lab ASC computing
- Ramp up to deliver user support for Sequoia
- Improve collaborative tools and self-help resources, particularly in support of ACES
- Further develop deep technical expertise for SNL participation in ACES platform support
- Develop a plan for addressing user support requirements for the NSCC
- Continue to leverage ITIL as a framework for improving the HPC OneStop service desk processes and practices
- Continue to deploy coordinated user training across SNL HPC OneStop user support sub-teams

Collaborations (WBS 1.5.5.3)

Program Support (LLNL)

Accomplishments in FY12

- Participated in planning with the HPC community for a possible future exascale program in partnership with the DOE/SC, including activities such as workshops and vendor meetings; the most significant accomplishment was development and completion of the extreme-scale FastForward procurement selection phase
- Continued management of the Sequoia contract and its associated R&D and design and engineering contracts with IBM; managed existing tri-lab contracts, including CCE and TLCC-related contract management
- Executed the FastForward contract process on behalf of NNSA and ASCR
- Completed the HUXTON and NUCLEI projects
- Supported PSAAP collaborations and PSAAP2 call, proposal review, and selection

Planned Activities in FY13:

- Continue FY12 procurement; continue contract management for TLCC, Sequoia, and FastForward; and continue extreme-scale computing planning
- Support workshops for DOE as needed

Program Support (LANL)

Accomplishments in FY12

- Hosted Predictive Science Panel (PSP) meeting at LANL in March 2012
- Coordinated and hosted ASC Principal Investigator meeting at Barksdale Air Force Base in February 2012
- Published ASC newsletter

Planned Activities in FY13:

- Host PSP Meeting
- Publish ASC newsletter

Program Support (SNL)

Accomplishments in FY12

- Supported two Qualification Alternatives to the Sandia Pulsed Reactor (QASPR) external review panel meetings
- Supported current PSAAP collaborations, and HQ and tri-lab community to review future PSAAP program proposals
- Provided support to the ASC Federal program office with emphasis on the Exascale Initiative, including JASONs Exascale Study
- Supported the 2011 Supercomputing Conference, PSP meetings, the ASC executive committee; quarterly meetings of the ASC executive committee, and managed the new SAIC contract to provide various administration support to HQ
- Designed and developed new SNL ASC Web page and content, with connection to NNSA ASC page

Planned Activities in FY13:

- Organize and host Fourth Predictive Engineering Science Panel meeting
- Support external review panel meetings for QASPR, the Engineering Sciences External Advisory Board, and the Computational Sciences External Advisory Board
- Support programmatic needs of the PSAAP program and the DOE Exascale Initiative
- Manage the SAIC contract to provide various administration support for HQ
- Support programmatic needs of NNSA tri-lab ASC program and ASC executive committee

Applications in Support of Manufacturing Production and Connectivity (Y-12)

Accomplishments in FY12

- Created the first quantitative, science-based models of Y-12's microwave operations. The models were developed to assist optimization of the new production microwave (PMW) unit that is presently being deployed by the Material Recovery and Recycle (MRR) organization to advance understanding of how geometries and materials determine microwave distributions, reflected power, and heating patterns, and permit MRR to rapidly and cheaply evaluate changes in geometry and process variables.
- Simulated radiation transport in selected regions of the Y-12 Uranium Processing Facility in support of the design for the Criticality Accident Alarm System (CAAS), which responds and alarms to a nuclear criticality. Effects of postulated nuclear criticality accidents were modeled to evaluate the adequacy of the facility design and verify whether design modifications are required to reduce radiation dose in particular regions. Investigated the application of advanced variance reduction techniques to MCNP Monte Carlo simulation models so that timely calculations can be run in production computing environments.
- Collaborated with the Kansas City Plant to optimize heat transfer within one of the reaction cells associated with development's direct electrolytic reduction/electro-refining (DER/ER) project. DER and ER are molten-salt-based electrochemical processes being developed to remove oxygen from metal oxide powders and then purify the resulting metals. Finite-element modeling performed by KCP engineers led to multiple changes in the geometry of the second-generation ER cell, as well as to a new DER cell design.
- Created a discrete-event simulation model of the Uranium Processing Facility with integrated graphics of the facility, equipment, and parts that can be used as the basis for a manufacturing process simulation to evaluate criticality response and evacuation routes from the facility.
- Utilized Y-12 and remote ASC cluster resources to solve production manufacturing problems in support of life extension programs.

Planned Activities in FY13:

- Benchmark GPU-based codes against similar codes ported to the APU-based architecture of the Penguin Altus 2A00 Llano platform at SNL
- Continue development and utilization of finite element models and the SALINAS code to estimate elastic or visco-elastic material properties in support of LEP
- Continue to evaluate new codes on the Y-12 cluster and utilize Y-12 and remote ASC cluster resources to solve production manufacturing problems
- Continue development of discrete-event nuclear facility simulation model
- Participate in National Security Enterprise ASC activities

System and Environment Administration and Operations (WBS 1.5.5.4)

System and Environment Administration and Operations (LLNL)

Accomplishments in FY12

- Implemented HPC Enclave Restricted Zone/Collaboration Zone security services
- Completed CRYPTO Card two-factor authentication deployment for HPC Enclave
- Enabled tri-lab inter-site authentication for Inter-Site High Performance Computing (IHPC)
- Deployed RSA one-time password two-factor authentication solution for secret national security information
- Migrated core security infrastructure components to RedHat Linux platform
- Integrated Sequoia
- Integrated Zin, Cab, rzmerl
- Removed and disposed of 42,000 tape cartridges

Planned Activities in FY13:

- Implement additional security and access management capabilities for ensuring the protection of resources in the HPC Enclave
- Consolidate infrastructure systems to a virtual machine environment
- Replace older network file system (NFS) scratch space servers
- Dismantle and dispose of uBlueGene/L, BlueGeneL, and Peloton-class systems
- Replace older Lustre servers
- Move Sequoia to classified environment
- Place Sequoia into general availability status
- Upgrade tape backup systems

System Administration and Storage (LANL)

Accomplishments in FY12

- Properly configured queues and scheduling policies; continued to conduct daily monitoring and problem resolution of user problems
- Provided data storage operations for GPFS, NFS, and archival storage (HPSS)
- Supported Cielo weapons science and integrated weapons applications across the tri-lab community
- Deployed newer Panasas file system technology to replace aging technology and increase production capacity computing scratch file system capacity
- Deployed PLFS on ASC production clusters

Planned Activities in FY13:

- Support HPC systems by conducting ongoing and daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Ensure workload is carried out by proper configuration of queues and scheduling policies plus daily monitoring and problem resolution relating to workloads running on HPC computing resources
- Enhance support for capacity computing utilization of GPUs
- Support requirements for transitioning unclassified HPC operations to the Turquoise network

Operations and Procurement Support (LANL)

Accomplishments in FY12

- Provided 24x7 operations and monitoring of HPC computing resources
- Provided hardware self-maintenance for ASC platforms
- Decommissioned Redtail and Viewmaster
- Implemented monitoring on new ASC systems and Lustre file system
- Procured HPC platforms, supporting hardware and software, and other products and services required by HPC

Planned Activities in FY13:

- Provide 24x7 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms
- Develop requirements for upgraded tools for monitoring HPC platforms and file systems
- Decommission Roadrunner/IBM system, Lobo, and Turing
- Provide technical and administrative support for procurement of HPC platforms, supporting hardware and software, and other products and services required by HPC

Requirements Planning (LANL)

Accomplishments in FY12

- Enhanced environmental data collection capabilities
- Designed and implemented infrastructure for enhanced system administrator troubleshooting tool
- Developed data feeds and filtering necessary for utilization monitoring

Planned Activities in FY13:

- Integrate Zenoss-based HPC monitoring into HPC data management infrastructure
- Implement requirements and design for upgrade of current monitoring system to latest open-source Zenoss version
- Design and implement interface for enhanced troubleshooting
- Develop and deploy proactive utilization monitoring within HPC monitoring

Computing Platform Integration and Deployment (LANL)

Accomplishments in FY12

- Installed Luna, the TLCC2 capacity system, and integrated the system into the LANL classified computing environment
- Installed Moonlight, a hybrid TLCC2 capacity cluster (GPGPU-enhanced TLCC2 cluster), into the LANL unclassified computing environment
- Completed the ASC Level 2 milestone review for Luna production readiness

Planned Activities in FY13:

- Upgrade the Luna system with two additional SUs
- Continue to operate Luna and the other capacity systems in both the classified and unclassified computing environments
- Complete the integration of Moonlight, a TLCC2 system with GPGPUS, into the LANL unclassified computing environment; work with the applications community for identification of additional software tools for the hybrid architecture
- Develop requirements for the next round of Commodity Technology Systems with tri-lab design team

Production Computing Services (SNL)

Accomplishments in FY12

- Integrated the TLCC2 platforms into production status
- Created robust TOSS installation and configuration process to ensure consistent, efficient upgrades on all HPC systems
- Instantiated a new Advanced Architectures Test Bed environment for use by SNL researchers and external collaborators
- Maintained production environments for ASC and nuclear weapons customers
- Supported ACES partnership in Cielo production through the CCC process; supported the Level 2 milestone for production capability

Planned Activities in FY13:

- Realign TLCC1 resources to serve foreign nationals (for example, PSAAP, employees)
- Continue operation support for Cielo and initial Trinity planning
- Continue to maintain production environments for ASC and nuclear weapons customers
- Manage operation of the DisCom WAN
- Move encrypted file system into production use within the NSCC and continue user support activities for remote users of these resources

Facilities, Network, and Power (WBS 1.5.5.5)

Facilities, Network, and Power (LLNL)

Accomplishments in FY12

- Installed B-453 petascale liquid cooling mechanical system; the new liquid cooling tertiary system is made up of pumps, heat exchangers, and piping
- Distributed B-453 petascale electrical power from the electrical room to the computer room
- Installed B-453 Sequoia utility infrastructure, including mechanical, electrical, and networking distribution to integrate Sequoia
- Upgraded B-451 Building Management System

Planned Activities in FY13:

- For the B-453 chilled water system chiller upgrade, replace small staging chiller with a larger in-line chiller to increase reliability and reduce nuisance staging of the chilled water system
- For the B-439 power and cooling improvement project, install electrical computational capacity by installing flexible distribution to the computer room (power distribution is antiquated and ease of expansion is limited), and move small staging chiller from B-453, utilizing B-439 for future cooling requirements
- For B-453 power realignment project, demolish and redistribute power for retired platforms (for example, BlueGene/L and Peloton) in preparation for future platforms and cluster installations, continue to enhance diagnostic and monitoring of IB fabrics on Sequoia and other IB attached Lustre file systems, deploy 40G/100G WAN connectivity in partnership with ESNet, and continue to analyze and evaluate emerging network technologies

Facilities, Networking, and Power (LANL)

Accomplishments in FY12

- Completed site preparation, analysis, and testing for SERF water integration
- Completed 60 percent of design phase for SCC upgrade in support of future ASC platforms
- Completed site preparation for ASC computing systems
- Maintained electrical and mechanical systems for ASC computing facilities

Planned Activities in FY13:

- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Switch to SERF water for cooling towers in SCC
- Start construction phase for SCC Trinity Infrastructure Project upgrade in support of future ASC platforms
- Integrate and deploy hardware and software in support of ASC computing on Turquoise network
- Provide ongoing operations and maintenance of HPC networking components

Facilities, Networking, and Power (SNL)

Accomplishments in FY12

- Retired and replaced Red Storm with NSCC2
- Upgraded DisCom network routers and encryptors
- Instantiated IHPC network
- Created remote graphics capability for SNL/CA applied to Special Compartmented Information (SCI) customers

Planned Activities in FY13:

- Continue growth of SCI resources and program support in the NSCC
- Plan expansion of NSCC facility and site additional SCI resources
- Consolidate, realign, or retire some HPC resources in B-880 Computing Annex to create space for new program needs
- Continue engagement with Solar Research partners at SNL on utilization of 2-MW solar panel array installed in FY12
- Explore alternative computing facilities (that is, Pods) for some users

Common Computing Environment (WBS 1.5.5.6)

System Software Deployment for Commodity Technology Systems

Accomplishments in FY12

- Released minor updates to TOSS (version 1.4-3 and 1.4-4), which included security updates and bug fixes, and supported upgrades of existing TOSS 1.x clusters to TOSS 2.0
- Developed/deployed a TOSS major update (version 2.0), which is based on RedHat's latest major release, RHEL 6
- Deployed the next generation of the ASC TLCC systems (TLCC2), which included software integration and testing for the tri-lab environment
- Deployed and tested TOSS extensions for GPU-enhanced clusters used for visualization and GPGPU workloads
- Developed HPL test plan for GPGPU; ran this on Moonlight (LANL cluster); discovered thermal issues, some power issues, and error-correcting code (ECC) memory issues that were subsequently addresses by Appro and NVIDIA

Planned Activities in FY13:

- Provide ongoing TOSS software development and support
- Provide support for GPU-enhanced clusters used for visualization and GPGPU workloads
- Develop/deploy TOSS 2.X (based on RHEL 6.X)

Programming Environment Development/Support for Tri-Lab Systems

Accomplishments in FY12

- Completed core CBTF infrastructure and transitioned O|SS performance collection tools to CBTF, ported the BlueGene/Q (Sequoia) environment, developed prototype system monitoring and administrative tools with CBTF
- Ported and tested debugging tools, including STAT and TotalView, for TLCC2; improved the STAT tool for Cielo (1.2.1 release) and ported to BlueGene /Q (initial scoping effort demonstrated potential for improving debugging capability in TotalView on templated C++ codes); the TotalView scalability contract is on track and producing improvements for tri-lab platforms, with scalable memory debugging milestone improved GUI performance among the deliverables
- With the University of New Mexico, progressed on merging of the lightweight infrastructure bootstrapping infrastructure and Multicast Reduction Network (MRNET) implementation into the main MRNET repository at University of Wisconsin
- For Open MPI development and support, developed new kernel-assisted shared-memory byte transfer layer (BTL), new open run-time environment (ORTE) routed component (*debruijn*) to address large-scale startup times in ORTE, mpimemu (a distributed memory usage tool that approximates memory resource utilization of MPI libraries), and various bug fixes to Open MPI to address specific tri-lab issues
- Continued activities with tool providers (O|SS, TAU, Valgrind) for various services, workshops, training, user support, fixes and problem resolution, and minor enhancements on user request

Planned Activities in FY13:

- Enhance the threading support and provide memory consumption tools in O|SS/CBTF framework; enhance support for I/O profiling and tracing
- Increase debugger tools capability to support increased scale of HPC systems, specifically TotalView and STAT
- Continue support contract efforts for open source tools (O|SS, TAU, HPCToolkit, MUST/Vampir, and Valgrind)
- Provide enhancements and bug fixes to Open MPI/MVAPICH based on tri-lab need, including, share MPI run-time parameters, assess MPI performance, identify best build and run parameters through studies, investigate impact of topology mapping, and provide results to end users

Resource Management Deployment and Reporting

Accomplishments in FY12

- Implemented plugins for task containment based on the new Linux cgroup capabilities available in TOSS 2.0
- Enhanced Workload Characterization (WC) Tool with improved custom reporting interface and with additional tools for configuration and administration
- Gathered requirements from NNSA/HQ and each laboratory to expand the reporting capabilities
- Improved WC Tool interface at each lab for data collection, storage, and reporting

Planned Activities in FY13:

- Contribute to SLURM development activities to provide functionality unique to the tri-labs
- Collectively address SLURM installation, configuration, management, testing, and diagnostic efforts that arise across the tri-Labs
- Evaluate SLURM's prospects for managing exascale clusters and increased resource integration
- Assess usage and work with ASC HQ to meet new and/or expanded HPC resource reporting requirements for WC Tool
- Improve WC Tool capabilities in reporting interface, administration and configuration tools, and metrics reporting

High Performance Computing Environment Integration for Tri-Lab Systems

Accomplishments in FY12

- Completed installation of 1-gigabyte (GB) networking infrastructure based on Interconnection Security Agreement design with the tri-lab, and upgraded 1-GB network encryptors to 10-GB encryptors.; Kerberos cross-realm trust was re-established between LLNL and LANL
- Developed detailed use case scenarios for IHPC to characterize the bounds of IHPC operation (enclave versus desktop versus enterprise network)
- Developed and tested Synchronized Account Request Automated Process (SARAPE) version 2; improved user interface with significantly enhanced documentation and context-sensitive help, ability to request resources at multiple sites in a single request, and to save/retrieve incomplete requests and ability to define multi-site services, as a combination of resources at multiple labs that define a single service
- Gathered initial requirements for SARAPE interface to local lab account processes

Planned Activities in FY13:

- Complete 10-GB networking infrastructure to the tri-lab
- Expand authentication capabilities to include Security Assertion Markup Language (SAML) for Web applications, including SARAPE; this expansion will allow IHPC to integrate with the work being done in the NNSA Network Vision project
- Investigate providing tri-lab Radius services to provide two-factor authentication for critical interface where two-factor authentication is mandatory, the use of a tri-lab shared file space, and unclassified disaster recovery leveraging IHPC resources
- Move SARAPE Version 2 into production for the tri-lab
- Increase trial/prototypes of collaborative tools (chat, JIRA, GitHub); continue maintenance of Gforge server for CCE project collaboration

Monitoring and Metrics Integration for Tri-Lab Systems

Accomplishments in FY12

- Established data services environment direction through interactions with W3C community, other DOE laboratory efforts, and prototypes by utilizing a representational state transfer software architectural approach and leveraging open source efforts for aggregation and Web-based data producer-consumer model tools
- Developed APIs to request data from sources, create a data response format, and map from common conventions to particular data source instances
- Developed common data naming conventions (Dictionary) for targeted use cases
- Deployed prototype services addressing defined use cases on systems

Planned Activities in FY13:

- Support development and use of a common set of Splunk tools across sites
- Implement generalized information sharing APIs within an HPC centric http-transport open source framework (LLNL's Lorenz or similar)
- Promote and collaborate on developing a commonly accepted API for http-transport data browsing/requests/response to be used in tool development
- Develop initial service-based tool prototype for high-fidelity characterization of resource usage (for example, CPU and memory utilization) for applications

File System Architecture and Integration

Accomplishments in FY12

- Demoed an app that writes each checkpoint as a single shared file using MPI-I/O and scalable/check-point restart (SCR) code via PLFS
- Added local storage support for checkpoint (memory and SSD/burst buffer), PLFS 2.1 (in TOSS 2.0) and PLFS 2.2
- Implemented and demoed at SC11 the PLFS I/O APIs to support moving data from/to burst buffer to/from parallel file system (Panasas's Panfs)

Planned Activities in FY13:

- Develop prototype to enable n–1 MPI-I/O apps to save checkpoints via PLFS/SCR
- Support testing and development of FastForward framework for asynchronous data transfer to PFS
- Investigate capability to compress and merge files when writing to PFS
- Demonstrate DAOS (object interface to Lustre)/POSIX
- Demonstrate asynchronous I/O and initial PLFS/DAOS VOL plugin and function shipping (VOL is an object interface at the bottom of HDF5)

V. ASC Level 1 and 2 Milestones

Table V-1. ASC Level 1 *Proposed* Milestones and Interfaces with Defense Programs Components from FY13–FY16

Milestone Title	Level	FY	Completion Date	Site(s)	Participating Program Offices
Advanced models to support initial conditions for boost (initial conditions 2)	1	FY14	TBD	LANL, LLNL	Science Campaigns ASC Campaign
Demonstrate predictive capability for weapon system response to neutrons in hostile radiation environment (maps to predictive capability framework peg post "Advanced Circuit Prediction and Integrated Response")	1	FY15	TBD	SNL	ASC Campaign Engineering Campaigns
Full-system integrated thermal safety assessment (maps to predictive capability framework peg post "3D Abnormal Assessment")	1	FY16	TBD	SNL	ASC Campaign Engineering Campaigns

ASC Level 2 Milestones for FY13

Table V-2. Quick Look: Level 2 Milestone Dependencies for FY13[6]

Milestone ID	Milestone Title	FY	Complete Date	Sub-Program	Site
TBD	Demonstrate Erosion with Auto-Contact Capability	FY13	12/31/12	IC	LLNL
TBD	Apply ASC Tools to Ignition-Relevant Experiments	FY13	6/30/13	IC	LLNL
TBD	Deliver Laser Ray Tracing Simulation Capability Required to Model Stockpile-Relevant Aboveground Experiments	FY13	9/30/13	IC	LLNL
TBD	Computation and Inter-Comparison of Four Distinct Equation of State Theory Predictions for the Lithium Deuteride Equation of State in an Interesting Regime	FY13	9/30/13	PEM	LLNL
TBD	Predictive Capability Assessment Project Results	FY13	9/30/13	V&V	LLNL
TBD	Report on Investigation of Input/Output Forwarding for Linux Clusters, as Part of Systems-Level Software at Scale	FY13	12/31/12	CSSE	LLNL
TBD	Sequoia File System Deployed to Classified Network	FY13	9/30/13	CSSE	LLNL
TBD	Early Users on Classified Sequoia Hardware	FY13	6/30/13	FOUS	LLNL
TBD	User Release in the Eulerian Application Project Codes of an Alternative Hydrodynamics Option	FY13	9/30/13	IC	LANL

[6] Factors such as FY13 Congressional Appropriations, NNSA/DP directives, and National Security considerations may necessitate a change in the current milestone set.

Milestone ID	Milestone Title	FY	Complete Date	Sub-Program	Site
TBD	Implement Improvements in the Lagrangian Application Codes for the FY14 Boost Initiation Peg Post	FY13	9/30/13	IC	LANL
TBD	Initial Multi-Phase Strength Model for a Single Phase Change	FY13	6/30/13	PEM	LANL
TBD	Development and Implementation of Initial Condition Models in Support of the FY14 Boost Initiation Peg Post	FY13	9/30/13	PEM	LANL
TBD	Quantification of Uncertainty Due to Numerical Errors and Approximations in Physics Calculations	FY13	6/30/13	V&V	LANL
TBD	Case Study of In Situ Data Analysis in ASC Integrated Codes	FY13	9/30/13	CSSE	LANL
TBD	Expanded Frequency Capability in EIGER for Electromagnetic Response Coupling	FY13	6/30/13	IC	SNL
TBD	New TRILINOS Solver Stack with Next-Generation High Performance Computing Platform Support Deployed for SIERRA/ThermalFluid	FY13	6/30/13	IC	SNL
TBD	Demonstration of Combined SIERRA/SolidMechanics and SIERRA/StructuralDynamics Code Capabilities with a Common Toolkit Underpinning	FY13	9/30/13	IC	SNL
TBD	Physics-Based Model for Lightning Arrestor Connector Performance	FY13	9/30/13	PEM	SNL

Milestone ID	Milestone Title	FY	Complete Date	Sub-Program	Site
TBD	Capability for Incorporating Micro-Scale Variability into Design and Analysis Models of 304L Stainless Steel Laser Welds	FY13	9/30/13	PEM	SNL
TBD	Sandia Physics and Engineering Models Roadmap	FY13	3/31/13	PEM	SNL
TBD	Predict the Aerodynamic Performance of a Full-Scale Stockpile Gravity Bomb with SIERRA Gas Dynamics Code	FY13	9/30/13	V&V	SNL
TBD	ASC V&V Project Credibility Assessment Using PCMM Methodology	FY13	9/30/13	V&V	SNL
TBD	Data Co-Processing for Extreme Scale Analysis	FY13	3/31/13	CSSE	SNL
TBD	SIROCCO File System Performance	FY13	6/30/13	CSSE	SNL
TBD	Demonstrate Effective SNL Access and Use of Sequoia System	FY13	6/30/13	FOUS	SNL
TBD	Deploy SNL Computing Resources onto the Inter-Site High Performance Computing Network	FY13	9/30/13	FOUS	SNL
TBD	Lustre File System Production on Cielo	FY13	3/31/13	CSSE	LANL SNL
TBD	Tri-Lab Data Backup and Recovery	FY13	9/30/13	FOUS	LLNL LANL SNL

Milestone (ID#): Report on Investigation of Input/Output Forwarding for Linux Clusters, as Part of Systems-Level Software at Scale		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 12/31/12		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: This milestone investigates I/O forwarding services for commodity Linux clusters to improve job launch scalability, reduce operating system jitter, and simplify system management. The investigation will discuss the motivation and requirements, review I/O forwarding architectures and implementations, and consider possible ways to deploy I/O forwarding on TLCC clusters as an alternative to direct-mounted NFS and Lustre file systems.		
Completion Criteria: Report to include results from investigation of I/O forwarding services for commodity Linux clusters.		
Customer: ASC/LLNL		
Milestone Certification Method: Professional report and hand-off to ASC program.		
Supporting Resources: Systems software expertise, test bed resources		

Milestone (ID#): Sequoia File System Deployed to Classified Network		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 9/30/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: This milestone represents the culmination of the procurement, development, integration, testing, and deployment of the 55-petabyte Sequoia file system (known as Grove. Grove is a Lustre file system that will serve as the primary parallel file system for the Sequoia platform, but will also be accessible from other LC machines. Completion of this milestone will require extensive integration and testing of Grove hardware and software components in the unclassified environment. The milestone will be complete when the Grove file system has been migrated to the classified environment and is available for use by production applications.		
Completion Criteria: Sequoia file system production operation in classified environment.		
Customer: ASC/LLNL		
Milestone Certification Method: Professional report and handoff to ASC program.		
Supporting Resources: Sequoia, systems, and file systems expertise		

Milestone (ID#): Early Users on Classified Sequoia Hardware		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 6/30/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The Sequoia system will provide approximately 20 peak petaFLOP/s of compute cycles to the center. The system was integrated from January through August 2012, with system acceptance scheduled for September 30, 2012. Unclassified science runs are scheduled to take place through January 2013. The system is expected to be moved to the classified network in mid-February 2013. The goal of this milestone is to have early users running their codes on Sequoia on the classified network.		
Completion Criteria: Racks have been assembled in the Terascale Simulation Facility (TSF), the system has been moved to the classified network, and at least one user has ported their code. A user will write a memo certifying that their code has run on Sequoia.		
Customer: ASC/LLNL		
Milestone Certification Method: Professional report and handoff to ASC program.		
Supporting Resources: Sequoia, code team participation		

Milestone (ID#): Case Study of In Situ Data Analysis in ASC Integrated Codes		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 9/30/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: The massive, complex datasets generated by scientific simulations on next-generation computer architectures at extreme scale pose new challenges for data management and scientific understanding. A critical element in meeting these new challenges is in situ data analysis where information is reduced and analyzed as close to the point of generation as possible, for example, at run time or as the data stream to longer term storage. This milestone will demonstrate in situ data analysis in xRAGE, an ASC Eulerian Radiation Hydrodynamics code. We will work with the Eulerian Applications Project to identify target problems that will exercise physics of interest to the program, and we will work closely with users to maximize the achieved usability and productivity improvements. We will investigate enhancements to the in situ process that enable and improve techniques of data reduction and data triage in addition to improving in situ performance by using portable acceleration technology. We will report on their viability and performance in xRAGE or another ASC code. These products will provide ASC IC developers with reference codes and tools that will benefit DSW by advancing the technology for predictive science at extreme scales.		
Completion Criteria: Demonstration of our in-situ visualization and analysis capability applied to xRAGE and/or other ASC codes with performance impacts documented.		
Customer: ASC IC development teams and DSW		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: System resources and participation of code teams and code users.		

Milestone (ID#): Data Co-Processing for Extreme Scale Analysis		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 3/31/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: ASC calculations produce complex datasets that are increasingly difficult to explore and understand using traditional post-processing workflows. To advance understanding of underlying physics, uncertainties, and results of ASC codes, SNL must gather as much relevant data as possible from large simulations. This drives SNL to couple data analysis and visualization capability with a running simulation, so that high fidelity data can be extracted and written to disk. This milestone evaluates two approaches for providing such a coupling: • In-situ processing provides "tightly-coupled" analysis capabilities through libraries linked directly with the simulation. SNL has collaborated on developing an in-situ capability designed for this purpose. • In-transit processing provides "loosely coupled" analysis capabilities by performing the analysis on separate processing resources. SNL provides this capability through a "data services" capability designed for this purpose.		
Completion Criteria: SNL will engineer, test, and evaluate customer-driven data operations on large-scale data created by a running simulation. The data operations will be performed by instrumented versions of both the in-situ and in-transit solutions, with the resulting performance data published and made available to the ASC community.		
Customer: ASC platform design team members and application developers.		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: Platform facilities, such as Cielo and associated test beds.		

Milestone (ID#): SIROCCO File System Performance		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 6/30/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Exascale computing will likely require fundamental changes in the storage and management of persistent data. Incremental advances in current capabilities are likely inadequate. This milestone will provide performance analysis of a revolutionary approach to persistent storage—one that uses smart storage servers with access to a variety of different local and remote media (for example, disk, NVRAM, memory, and tape) and are pervasive throughout the computing platform. Storage servers have the ability to directly handle I/O requests, initiate third party transfers, or replicate the data as needed. Results will come from implementations for the Cielo system.		
Completion Criteria: Completion of the Sirocco performance analysis.		
Customer: ASC platform design team members and application developers.		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: Cielo and associated test beds		

Milestone (ID#): Demonstrate Effective SNL Access and Use of Sequoia System		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 6/30/13		
ASC nWBS Subprogram: FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Determine, document, and publish the recommended access methods for SNL HPC customers (and local user support personnel) to the Sequoia computer system hosted at LLNL. Prepare guidelines for job submission, data transfer to and from LLNL, trouble ticket generation, and coordination with LLNL Help Desk personnel. Assist in debugging code or scripting as necessary to ensure success for the initial SNL users, especially those requesting CCC access.		
Completion Criteria: Report or presentation providing evidence of proper accounting processes, job submission, data transfer, visualization of input and/or output for production-type jobs.		
Customer: ASC and DSW		
Milestone Certification Method: Testimony (email, presentations, graphics results) of success on Sequoia runs for SNL users.		
Supporting Resources: Long-haul communication lines between labs.		

Milestone (ID#): Deploy SNL Computing Resources onto the Inter-Site High Performance Computing Network		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 9/30/13		
ASC nWBS Subprogram: FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Identify cluster systems and additional hardware that can move from the Sandia restricted network address space into the IHPC address space to provide enhanced resources and broader access for SNL external and internal customers. Likely candidates are the Glory TLCC1 HPC cluster and some Lustre file system servers. Prepare the IHPC environment for these systems and move them no later than June 30, 2013.		
Completion Criteria: The candidate equipment will be operational and effective use, including external to SNL, shall be demonstrated.		
Customer: ASC environment customers, including external partners/collaborators.		
Milestone Certification Method: Final documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a customer is documented.		
Supporting Resources: SNL FOUS operations teams and CCE IHPC project.		

Milestone (ID#): Lustre File System Production on Cielo		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 3/31/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: The Lustre file system on Cielo is providing reliability and performance needed for full production of the Cielo platform. Application users can get their work done on Cielo without major file system issues. Requirements include: file system deployment is complete, acceptance testing has been completed, reliability and performance of the file system has been demonstrated by application program, all hardware for I/O is in operation, and Lustre is effectively deployed into the full Cielo environment, including purging, file transfer agent access, external node access, and monitoring and management. A milestone review has been completed and the results documented.		
Completion Criteria: Follows the ASC Level 2 Milestone criteria for capability platforms but specifically with regards to the Lustre file system on Cielo: that the file system has been successfully demonstrated for production capability-class simulations. The file system requirements are specifically listed in the usage model for Cielo, which defines that the Cielo file system has demonstrated an acceptable production user environment with all the associated support, testing, reliability, performance, and applications use of the system.		
Customer: NNSA/ASC HQ, ASC program managers responsible for CCCs, SSP, and ASC tri-lab CCC users.		
Milestone Certification Method: A program review is conducted, requiring all three laboratories to participate in the completion of the review. The review panel will include a mix of systems personnel and end users. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: ACES personnel, Cielo, and associated hardware test beds		

Milestone (ID#): Tri-Lab Data Backup and Recovery		
Level: 2	**Fiscal Year**: FY13	**DOE Area/Campaign**: ASC
Completion Date: 9/30/13		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL, LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: A code team from each laboratory will recover an agreed upon application at a remote lab using their backup copy at the remote lab. Once the software application has been recovered, the code team will demonstrate that the application was recovered successfully.		
Completion Criteria: A code team from LANL, LLNL, and SNL has successfully recovered and demonstrated that the chosen application is working at a remote lab using the off-site backup copy at the remote lab.		
Customer: ASC code teams		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: HPSS and IC, V&V, or PEM project team member support from each site		

VI. ASC Performance-Based Initiatives for FY13

The contractor's *Performance Evaluation Plan* contains multi-site targets that can be identified by the associate deputy administrator as base or stretch goals.

There are no multi-site targets for ASC in FY13.

VII. Performance Measures

Table VII-1. ASC Campaign Annual Performance Results (R) and Targets (T)

REDUCED RELIANCE ON CALIBRATION: The cumulative percentage reduction in the use of calibration "knobs" to successfully simulate nuclear weapons performance (Long-term Outcome)

Reporting process: Labs will assess and report within two weeks of the end of Q3, progress toward Level-1 milestones for knob replacement and supporting Level-2 milestones. A briefing will be provided to the Predictive Science Panel at each meeting as a report on progress toward replacing the knobs. The PSP may provide advisory comments regarding the reported progress, but will not evaluate the progress.

At LLNL, progress on this metric was initially measured by assigning 25% completion to each of the original 4 knobs. In 2010 the first of those knobs was complete. However, these 4 knobs have now been replaced by 8 PCF pegposts, the second of which will complete in 2012. Partial progress toward the next major pegpost can be achieved by successful completion of milestones related to predictive capability. However, due to delays in NIF weapons science experiments as well as delays in time projection chamber and chi-nu nuclear data experiments we are not able to report being fully at target in FY12.

LANL: Progress on this metric will be measured by completion of the number of Level 1 milestone that directly address knob replacement. Completing one knob-replacement Level 1 milestone accumulates 100%/N reduced reliance on calibration. Partial progress toward the next knob-replacement Level 1 milestone is achieved by successful completion of predictive capability Level 2 milestones. Cumulative percentage is reported for July 1, 2005 though June 30, 2012 (FY05Q4 through FY12Q3). In 2011 we rebaselined this metric over 8 major pegposts of the Predictive Capability Framework, 2 pegposts per knob (N=8-1) (i.e., the number of PCF Pegposts expected to be L1s since reconfiguration in FY11)

Evidence (Type): NA-10 Milestone Reporting Tool (MRT) reports (Original Documents)

	Q3 FY07	Q3 FY08	Q3 FY09	Q3 FY10	Q3 FY11	Q3 FY12	Q3 FY13	Q3 FY14	Q3 FY15	Q3 FY16	Q3 FY17	Q3 FY18
TARGETS	8%	16%	25%	30%	35%	40%	45%	50%	55%	60%	65%	70%
LLNL	12%	18%	24%	33%	35%	37%						
LANL	12%	18%	25%	33%	35%	37%						
AVERAGE	12%	18%	25%	33%	35%	37%						

VIII. Budget

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.1 Integrated Codes					
Engineering and Physics Integrated Codes		73.441	0.000	8.402	81.843
	LLNL	26.078			26.078
	LANL	23.808			23.808
	SNL	23.555			23.555
	Other			8.402	8.402
Specialized Codes and Libraries		27.290	0.000	0.000	27.290
	LLNL	8.845			8.845
	LANL	10.145			10.145
	SNL	7.300			7.300
	Other	1.000			1.000
Applications and Algorithms Research		15.733	0.000	0.000	15.733
	LLNL	5.456			5.456
	LANL	5.777			5.777
	SNL	4.000			4.000
	Other	0.500			0.500
Applications Research for Next-Gen Platforms		12.546	0.000	11.750	24.296
	LLNL	2.948			2.948
	LANL	5.248			5.248
	SNL	4.350			4.350
	Other			11.750	11.750

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.2 Physics and Engineering Models					
High Explosive		5.553	0.000	2.725	8.278
	LLNL	2.107			2.107
	LANL	3.446			3.446
	SNL				0.000
	Other			2.725	2.725
Equation of State		7.397	0.000	0.000	7.397
	LLNL	3.132			3.132
	LANL	3.665			3.665
	SNL	0.600			0.600
	Other				0.000
Nuclear Properties		5.008	0.000	0.000	5.008
	LLNL	1.671			1.671
	LANL	3.337			3.337
	SNL				0.000
	Other				0.000
Plasma and Radiative Properties		5.166	0.000	0.000	5.166
	LLNL	1.671			1.671
	LANL	3.495			3.495
	SNL				0.000
	Other				0.000
Advanced Hydrodynamics		6.181	0.000	0.000	6.181
	LLNL	1.671			1.671
	LANL	4.510			4.510
	SNL				0.000
	Other				0.000
Material Strength and Damage		13.102	0.000	0.000	13.102
	LLNL	3.963			3.963
	LANL	4.189			4.189
	SNL	4.950			4.950
	Other				0.000
Forensics and Cross-cutting Initiatives		8.737	0.000	0.000	8.737
	LLNL	3.296			3.296
	LANL	3.387			3.387
	SNL	2.054			2.054
	Other				0.000
Thermal and Fluid Response		5.100	0.000	0.000	5.100
	LLNL				0.000
	LANL				0.000
	SNL	5.100			5.100
	Other				0.000
Aerodynamics and Vibration		3.950	0.000	0.000	3.950
	LLNL				0.000
	LANL				0.000
	SNL	3.950			3.950
	Other				0.000
Radiation and Electrical Response		4.518	0.000	0.000	4.518
	LLNL				0.000
	LANL				0.000
	SNL	4.518			4.518
	Other				0.000
Russian Programs		2.462	0.000	0.010	2.472
	LLNL	0.303			0.303
	LANL	0.659			0.659
	SNL	0.375			0.375
	Other	1.125		0.010	1.135

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.3 Verification and Validation					
V&V Methods		22.276	0.000	0.000	22.276
	LLNL	5.260			5.260
	LANL	8.956			8.956
	SNL	5.400			5.400
	Other	2.660			2.660
V&V Assessments		25.555	0.000	0.000	25.555
	LLNL	4.594			4.594
	LANL	5.261			5.261
	SNL	15.700			15.700
	Other				0.000
Data Validation, Archiving, SQA and Training		9.701	0.000	0.000	9.701
	LLNL	4.634			4.634
	LANL	3.717			3.717
	SNL	1.350			1.350
	Other				0.000

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.4 Computational Systems and Software Environment					
Commodity Technology Systems		5.120	7.229	0.000	12.349
formerly Production Systems	LLNL	0.912	4.329		5.241
	LANL	1.155	2.900		4.055
	SNL				0.000
	Other	3.053			3.053
Advanced Technology		3.884	60.200	3.000	67.084
formerly Advanced Systems	LLNL	0.084	37.200		37.284
	LANL				0.000
	SNL	3.800		3.000	6.800
	Other		23.000		23.000
System Software and Tools		20.606	0.000	0.000	20.606
	LLNL	5.174			5.174
	LANL	8.032			8.032
	SNL	7.400			7.400
	Other				0.000
Input/Output, Storage, and Networking		13.245	0.000	0.000	13.245
	LLNL	7.476			7.476
	LANL	1.019			1.019
	SNL	4.750			4.750
	Other				0.000
Post-processing Environments		10.554	0.000	0.000	10.554
	LLNL	1.833			1.833
	LANL	5.021			5.021
	SNL	3.700			3.700
	Other				0.000
Next-Generation Computing Technologies		0.000	0.000	12.500	12.500
	LLNL			12.500	12.500
	LANL				0.000
	SNL				0.000
	Other				0.000

Product / Project WBS	Lab	PEOPLE ($Ms)	INFOSTRUCTURE ($Ms)	CONTRACTS ($Ms)	Total ($Ms)
1.5.5 Facility Operations and User Support					
User Support Services		14.463	0.000	23.472	37.935
	LLNL	3.074		1.083	4.157
	LANL	4.929			4.929
	SNL	2.200			2.200
	Other	4.260		22.389	26.649
Collaborations		3.874	0.000	8.128	12.002
	LLNL	1.632		0.050	1.682
	LANL	1.802		0.465	2.267
	SNL	0.440		1.037	1.477
	Other			6.576	6.576
System and Environment Administration and Operations		36.588	54.378	0.000	90.966
	LLNL	10.900	30.000		40.900
	LANL	13.519	18.747		32.266
	SNL	12.169	5.631		17.800
	Other				0.000
Facilities, Network and Power		8.259	26.540	1.450	36.249
	LLNL	3.759	10.000		13.759
	LANL	3.000	11.040		14.040
	SNL	1.500	5.500	1.200	8.200
	Other			0.250	0.250
Common Computing Environment		9.112	0.000	0.795	9.907
	LLNL	3.068		0.795	3.863
	LANL	3.307			3.307
	SNL	2.737			2.737
	Other				0.000

	PEOPLE	INFOSTRUCTURE	CONTRACTS	Total
Integrated Codes	129.010	0.000	20.152	149.162
Physics and Engineering Models	67.174	0.000	2.735	69.909
Verification and Validation	57.532	0.000	0.000	57.532
Computational Systems and Software Engineering	53.409	67.429	15.500	136.338
Facility Operations and User Support	72.296	80.918	33.845	187.059
Total	379.421	148.347	72.232	600.000
	63%	25%	12%	600.000

Month	Monthly Cost	Cumulative Cost
Oct		
Nov		
Dec		
Jan		
Feb		
Mar		
Apr		
May		
Jun		
Jul		
Aug		
Sep		

Appendix A. Glossary

2D	Two Dimensional
3D	Three Dimensional
ACES	New Mexico Alliance for Computing at Extreme Scale
ACSLS	Oracle's Automated Cartridge System Library Software
ADEPT	Applications Development Environment and Performance Team
API	Application Programming Interface
APU	Accelerated Processing Unit
ASC	Advanced Simulation and Computing
ASCI	Accelerated Strategic Computing Initiative
ASCR	Advanced Scientific Computing Research (DOE/SC)
BTL	Byte Transfer Layer
CBTF	Component-Based Tool Framework
CCC	Capability Computing Campaign
CCE	Common Computing Environment
CD	Critical Decision
CPU	Central Processing Unit
CRASH	Center for Radiative Shock Hydrodynamics, University of Michigan
CSSE	Computational Systems and Software Environment
DAG	Directed Acyclic Graph
DDS	Distributed Data Services
DOE	Department of Energy
DSW	Directed Stockpile Work
ECC	Error-Correcting Code
FOUS	Facility Operations and User Support
GB	Gigabytes
GPFS	Global Parallel File System
GPGPU	General-Purpose Graphics Processing Units
GPU	Graphics Processing Unit

HPC	High Performance Computing
HPCIC	High Performance Computing Innovation Center
HPSS	High-Performance Storage System
HQ	ASC Headquarters
I/O	Input/Output
IC	Integrated Codes
IHPC	Inter-Site High Performance Computing
ITIL	Information Technology Infrastructure Library
LC	Livermore Computing
LEP	Life Extension Program
LLNL	Lawrence Livermore National Laboratory
MIC	Many Integrated Core
mini-app	Mini Application
MPI	Message Passing Interface
MRNET	Multicast Reduction Network
NAS	Network-Attached Storage
NFS	Network File System
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NSCC	National Security Computing Center
NUMA	Non-Uniform Memory Access
nWBS	National Work Breakdown Structure
O|SS	Open|SpeedShop
ORTE	Open Run-Time Environment
PCF	Predictive Capability Framework
PDE	Partial Differential Equation
PEM	Physics and Engineering Models
PFTP	Parallel File Transfer Protocol
PLFS	Parallel Log File System
PRISM	(Center for) Prediction of Reliability, Integrity, and Survivability of Microsystems, Purdue University
PSAAP	Predictive Science Academic Alliance Program

PSI	Parallel Storage Interface
PSP	Predictive Science Panel
QMU	Quantification of Margins and Uncertainties
QASPR	Qualification Alternatives to the Sandia Pulsed Reactor
R&D	Research and Development
RAIT	Redundant Array of Independent Tapes
RANS	Reynolds-Averaged Navier Stokes
RAS	Reliability, Availability, Scalability
RHEL	Red Hat Enterprise Linux
SAN	Storage Area Network
SAP	Scalable Applications Project
SAML	Security Assertion Markup Language
SARAPE	Synchronized Account Request Automated Process
SC	Department of Energy's Office of Science
SCC	Strategic Computing Complex
SCI	Special Compartmented Information
SCR	Scalable/Check-Point Restart
SERF	Sanitary Effluent Reclamation Facility
SFI	Significant Finding Investigation
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SQA	Software Quality Assurance
SSD	Solid-State Drive
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
STAT	Stack Trace Analysis Tool
SU	Scalable Unit(s)
TLCC	Tri-Lab Linux Capacity Cluster
TOSS	Tripod Operating System Software
TSF	Terascale Simulation Facility
UQ	Uncertainty Quantification
V&V	Verification and Validation

VM	Virtual Machine
VTK	Visualization Toolkit
WAN	Wide Area Network
WBS	Work Breakdown Structure
WC Tool	Workload Characterization Tool

Appendix B. Points of Contact

WBS	Title	Contact
1.5.4	Computational Systems and Software Environment	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.8	Commodity Technology Systems	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.3	Advanced Technology Systems	Bronis de Supinski, LLNL, 925-422-1062, desupinski1@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Doug Doerfler, SNL, 505-844-9528, dwdoerf@sandia.gov
1.5.4.9	Next-Generation Computing Technologies	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov Jim Ang, SNL, 505-844-0068, jaang@sandia.gov
1.5.4.4	System Software and Tools	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov Sue Kelly, SNL, 505-845-9770, smkelly@sandia.gov
1.5.4.5	I/O, Storage Systems, and Networking	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov Ron Oldfield, SNL, 505-284-9153, raoldfi@sandia.gov
1.5.4.6	Post-Processing Environments	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov David Rogers, SNL, 505-844-5323, dhroger@sandia.gov
1.5.5	Facility Operations and User Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Tom Klitsner, SNL, 505-844-1901, tklitsn@sandia.gov

WBS	Title	Contact
1.5.5.2	User Support Services	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov
1.5.5.3	Collaborations	Brian Carnes, LLNL, 925-423-9181, carnes1@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Justine Johannes, SNL, 505-844-1994, jejohan@sandia.gov
1.5.5.4	System and Environment Administration and Operations	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov
1.5.5.5	Facilities, Network, and Power	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov
1.5.5.6	Common Computing Environment	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Judy Sturtevant, SNL, 505-845-9448, jesturt@sandia.gov
Caltech, Center for Simulating Dynamic Response of Materials		Michael Ortiz, 626-395-4530, Ortiz@aeor.caltech.edu Mark Stalzer, 626-395-2521, stalzer@caltech.edu Susan Powell, 626-395-2909, spowell@cacr.caltech.edu
Purdue University, Center for Prediction of Reliability, Integrity and Survivability of Microsystems		Jayathi Murthy, 765-494-5701, jmurthy@ecn.purdue.edu Matt Potrawski, 765-496-9301, mpotraws@purdue.edu
Stanford, Center for Integrated Turbulence Simulation		Parviz Moin, 650-723-9713, moin@stanford.edu Gianluca Iaccarino, 650-723-9599, iops@stanford.edu Deb Michael, 650-725-2077, debmich@stanford.edu
University of Michigan, Center for Radiative Shock Hydrodynamics		Paul Drake, 734-763-4072, rpdrake@umich.edu James Holloway, 734-936-3126, hagar@umich.edu
University of Texas, at Austin, Center for Predictive Engineering and Computational Sciences		Bob Moser, 512-471-0093, rmoser@mail.utexas.edu Chris Simmons, 512-232-2881, csim@colab.ices.utexas.edu

Appendix C. Academic Alliance Centers

California Institute of Technology

The Center for the Predictive Modeling and Simulation of High-Energy Density Dynamic Response of Materials

Accomplishments in FY12

- Upgraded the Small Particle Hypervelocity Impact Range (SPHIR) facility, which now houses a broad array of diagnostics including profilometry, 3D debris shadowgraph, impact-flash spectrometry, and 3D debris soft capture
- Completed 45 percent of the experimental campaign, Nylon 6/6 impactors and 6061-T6 Aluminum target plates at velocities in the range of 5–6 Km/s
- Demonstrated extension of EOS and viscosity laws of Nylon 6/6 and aluminum to the range of temperatures and pressures achieved during hypervelocity impact
- Developed and demonstrated multimaterial, fully 3D Eulerian capability for the simulation of hypervelocity impact, including strength and plasticity
- Performed extensive floating-point operation performance optimization of the OTM Lagrangian solver, resulting in a factor of five speed-up in single-node performance
- Achieved multi-teraFLOP/s performance by the OTM Lagrangian solver on Hera with excellent scalability up to 7612 cores
- Developed an optimal UQ protocol based on legacy data and implemented it into the hydra UQ pipeline and surrogate model demonstration
- Completed 75percent of the annual UQ campaign, bringing together the Eulerian and Lagrangian solvers, the extended material models, the Legacy-OUQ protocol, and the hypervelocity legacy SPHIR data

Planned Activities in FY13:

- Perform oblique hypervelocity SPHIR tests of single and double aluminum/steel biomaterial plates and nylon projectiles, 5–10 km/s velocity range, and compile perforation, impact flash spectroscopy, and 3D debris shadowgraphs and capture data
- Simulate Lagrangian and Eulerian SPHIR hypervelocity tests
- Use Lagrangian and Eulerian solvers and hydra UQ pipeline for optimal-UQ legacy-data-based analysis of hypervelocity impact experiment
- Demonstrate petascale performance of the Lagrangian and Eulerian solvers on Sequoia
- Validate multiscale iron model using in-house EBSD and DIC data along with Omega and NIF data (leveraged from HEDLP grant)

Purdue

Center for Prediction of Reliability, Integrity, and Survivability of Microsystems

Accomplishments in FY12

- Developed and implemented an integrated UQ framework based on Bayes networks for complex multiscale multiphysics applications in MEMS, accounting for heterogeneity in length and time scales and in device geometry.
- Made a successful first-pass at Bayes network analysis of pull-in, pull-out, dielectric charging and creep in MEMS employing heterogeneous data on frog-leg, cantilever and Center for Prediction of Reliability, Integrity, and Survivability of Microsystems (PRISM) devices.
- Completed development of main physical models, numerics and software and their integration, as well as testing and verification.
- Developed first-of-their kind models for contact and dielectric charging integrating atomistic physics with device-level models, and accounting for random inputs.
- Made first-of-their-kind integrated predictions of MEMS behavior, spanning length and time scales from density functional theory (DFT) to the device scale; made probabilistic predictions of device behavior in single and multicycle contact and creep.

Planned Activities in FY13:

- Continue experiments on multi-cycle contact, pull-in, pull-out, and creep using the PRISM and cantilever switches
- Refine creep model using new creep data for cantilevers; employ the Bayesian framework for validation and calibration
- Refine multiple trap depth dielectric charging model using new cantilever data for sustained contact; employ the Bayesian framework for calibration and validation
- Implement, test, and verify COMET algorithm for acceleration of ES-BGSK computations; implement and demonstrate scaling on 100–1000 processors; integrate, test, and verify immersed boundary method with COMET.
- Continue to perform and refine probabilistic integrated simulations of 1) multi cycle contact simulations of PRISM switch with coupled fluid-structure-electrostatics interactions, contact and dielectric charging, and prediction of dielectric charging rate; 2) sustained long term contact simulations in the presence of dielectric charging, and prediction of pull-out voltage after sustained contact; and 3) long-term creep simulations and prediction of gap versus time; perform Bayesian UQ analysis for all these simulations, and refine simulations as necessary

Stanford University

The Center for Predictive Simulations of Multi-Physics Flow Phenomena with Application to Integrated Hypersonic Systems

Accomplishments in FY12

- Defined a multiple-gate QMU/UQ analysis framework; verified that the concept of separating gates can be used to reduce the dimension of the stochastic space for the specific problem of uncertain chemical reaction rates
- Developed approach to quantify epistemic uncertainties in the turbulence model that incorporates physical constraints and goes beyond parametric uncertainties.
- Developed multi-physics adjoint solver (used for sensitivities and verification) for turbulent flow with tabulated chemistry
- Completed detailed comparison of shock/boundary-layer simulations (including both Reynolds-Averaged Navier Stokes (RANS) and large eddy simulation) and in-house experiments
- Demonstrated levels of performance equal to those of hand-tuned code for four different applications written in the Liszt Domain Specific language, while retaining automatic portability to distributed- and shared-memory machines, as well as GPUs

Planned Activities in FY13:

- Develop and test a methodology to investigate uncertainties in the kinetic reaction rates in full-system simulation of the Hyshot vehicle
- Complete comprehensive experimental database on supersonic combustion, including novel Toluene-PLIF and Tunable Diode Laser Spectroscopy measurements; complete validation of the flamelet-based combustion model
- Demonstrate full QMU/UQ methodology, including aleatoric, epistemic, and numerical error estimation on two sub-system problems
- Apply RANS solver written in the domain-specific language Liszt to UQ problems on GPUs and heterogeneous systems
- Complete the study of the sensitivity of shock/boundary-layer interaction under geometrical uncertainty using RANS, LES, and experiments
- Develop revised chemical mechanism for JP-7 surrogate fuel with measured uncertainties; quantify how uncertainties in the reaction rates propagate to the flow field in the HIFiRE-2 scramjet
- Develop methodology for UQ of radiation heat transfer in supersonic reacting flows
- Apply advanced laser diagnostics (developed within the PSAAP Center) to new full-system experiments at DLR in Germany and to the Stanford model combustor

University of Michigan

Accomplishments in FY12

- Enabled the CRASH laser package to trace rays in 3D and deposit energy in 2D or 3D
- Completed a UQ run set using the CRASH laser package and several other run sets
- Completed a UQ analysis that combined models of varying fidelity
- Performed simulations of the (3D) CRASH year-5 experiment
- Performed test experiments with nozzles and elliptical tubes

Planned Activities in FY13:

- Assess the predictive capability of the CRASH code by providing probability distributions for the year-5 experiment in advance
- Evaluate and report the successes and failures of these predictions after the experiments
- Extend the studies using CRASH and PDT to assess the errors associated with using multi group diffusion for the radiation energy transport
- Continue work on improved Krylov preconditioning in the CRASH code
- Perform an Omega day of Year-5 experiments and analyze the results

University of Texas

The Center for Predictive Engineering and Computational Sciences

Accomplishments in FY12

- Integrated FIN-S with updated physics models in turbulence, transport, and chemistry as guided by validation
- Completed development of an arc-jet simulator
- Completed calibration and initial validation in all single physics domains
- Developed new predictive validation framework including predictive assessment
- Explored adjoint-enhanced forward propagation of uncertainty

Planned Activities in FY13:

- Continue development of coupled code for modeling Arc Jet flows
- Develop spectral radiation code using quantum calculations for transition probabilities
- Calibrate and validate RANS models and their inadequacy models of boundary layers in hypersonic flows
- Continue development of the predictive validation framework
- Complete propagation of uncertainty in the full system simulation using validated single physics and validated multi physics

Appendix D.
ASC Obligation/Spend Plan

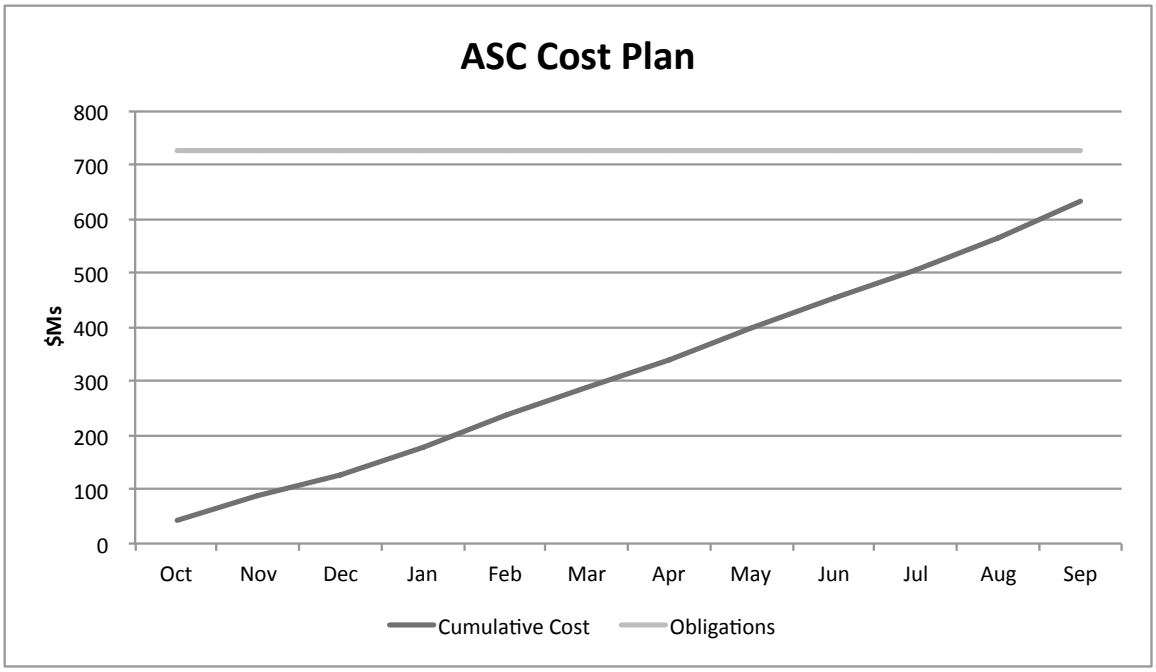

Figure D-1. ASC obligation/spend plan for FY13.


Advanced
Simulation &
Computing™